



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN

Economics

Ciclo XXXIV

Settore Concorsuale: 13/A1 – ECONOMIA POLITICA

Settore Scientifico Disciplinare: SECS-P/01 – ECONOMIA POLITICA

**ESSAYS IN MACROECONOMICS:
insurance, market power, and inequality**

Presentata da: [Oliviero Pallanch](#)

Coordinatore Dottorato

Andrea Mattozzi

Supervisore

Antonio Minniti

Esame finale anno 2023

Essays in Macroeconomics:
insurance, market power, and inequality

Oliviero Pallanch

December 2023

Contents

1	Weather shocks and coping strategies in Ethiopia	11
1.1	Introduction	11
1.1.1	Overview	11
1.1.2	Literature Review	12
1.1.3	Contributions	16
1.2	Data	18
1.2.1	Households Survey	18
1.2.2	Weather data	25
1.2.3	The matching procedure	36
1.3	Theory and Empirical Specification	37
1.3.1	Theoretical framework	37
1.3.2	Empirical framework	39
1.4	Results	41
1.4.1	Shocks and outcome	41
1.4.2	Robustness and Heterogeneity	44
1.4.3	Coping strategies	48
1.5	Conclusions	50
1.A	Robustness	52
1.A.1	Robustness and heterogeneity	52
1.A.2	Spatial features	52
1.B	Other specifications	53
1.B.1	Shocks and outcomes	53
1.B.2	Coping strategies	54
1.B.3	First differences	54
	References	58
2	Trends of concentration in selected OECD countries	63
2.1	Introduction	63
2.2	Literature Review	64
2.2.1	Concentration	64
2.2.2	Entrenchment	67
2.2.3	Explanations of recent trends in competition proxies	68
2.2.4	Contribution of this report to the literature	69
2.3	Data	71
2.3.1	Overview	71
2.3.2	Production data	72

2.3.3	Trade data	74
2.3.4	Firm-level financial data	76
2.3.5	Firm ownership information	77
2.4	Methodology	78
2.4.1	Overview	78
2.4.2	Taxonomy of industries	79
2.4.3	Concentration	82
2.4.4	Entrenchment	88
2.5	Trends	90
2.5.1	Concentration	90
2.5.2	Entrenchment	100
2.6	Conclusions	105
2.A	Data	108
2.A.1	Production data	108
2.A.2	Firm level financial data	116
2.A.3	Ownership data	118
2.A.4	Sample for concentration and entrenchment	120
2.B	Methodology	122
2.B.1	Trends	123
	References	130
3	A note on technological change and inequality	133
3.1	Introduction	133
3.1.1	Overview	133
3.1.2	Literature review	137
3.2	Theoretical framework	139
3.2.1	A static model	139
3.2.2	A dynamic model	145
3.3	Conclusions	150
3.A	Entrepreneurial Technology	151
3.B	Constrained Profit Maximization	151
3.C	Occupational Choice	154
	References	158

List of Figures

1.1	Survey Enumeration Areas distribution over Ethiopia	19
1.2	Survey timing across waves	20
1.3	Crop Calendar in Ethiopia	27
1.4	Agro-ecological zones as per survey	28
1.5	Distributions of Shocks across months of 6-months SPEI	31
1.6	6-months SPEI values in December.	32
1.7	March-September long-run precipitation level (mm)	34
1.8	Precipitation anomalies and long-run levels	35
1.9	Matching procedure	36
2.1	Apportioning for Production Data	73
2.2	Trade data apportioning	75
2.3	Example of the apportioning technique of Business Group activities	86
2.4	Concentration levels across geographical buckets	91
2.5	Concentration cumulative changes across geographical buckets	92
2.6	Trade adjustments, industries competing domestically	94
2.7	Trade adjustments, industries competing at the European level	95
2.8	Trade adjustments, industries competing at the global level	96
2.9	Concentration levels across selected quantiles	97
2.10	Cumulative changes across selected quantiles	98
2.11	Leadership ratio across geographical buckets (1st over 2nd), unweighted	99
2.12	Leadership ratio: alternative measures	100
2.13	Entrenchment across geographical buckets	101
2.14	Entrenchment for different time horizons	102
2.15	Share of markets by new top4 firms within each geographical bucket	103
2.16	Number of years that a firm stays in the top 4 and top1	104
2.17	Example of correspondence	114
2.18	Filling Backward Example	115
2.19	Apportioning procedure example	116
2.20	Concentration levels across geographical buckets (weighted)	123
2.21	Concentration cumulative changes across geographical buckets (weighted)	123
2.22	Concentration level treating all industries as European (unweighted)	124
2.23	Concentration cumulative changes treating all industries as European(unweighted)	124
2.24	Number of distinct industries per decile of concentration	125
2.25	Industries in the top decile of the concentration distribution	126
2.26	Leadership ratio across geographical buckets (1st vs. 2nd), weighted	127
2.27	Changes in the top4 ranking with no new firms in the top4, domestic bucket	127

2.28	Changes in the top 4 ranking with one new firm in the top 4, domestic bucket	128
2.29	Changes in the top4 ranking in markets with no entry, European bucket	129
3.1	Evolution of pre-tax income shares earned by the top earners	135
3.2	Wealth (net worth) shares owned by the richest fraction of the wealth distribution	136
3.3	Labor productivity: frontier vs.laggard firms	136
3.4	Characterisation of equilibrium at two different technology level, M_{low} and M_{high} .	140
3.5	Labour excess demand for different technological levels	141
3.6	Equilibrium wage and occupational choice cutoff at different levels of technology M	142
3.7	Top income shares at different levels of technology M	144
3.8	Simple counterfactual exercise	145
3.9	Preliminary computations	149
3.10	Top to bottom: decreasing credit frictions (increasing λ)	153
3.11	Occupational Choice, given wealth.	155
3.12	Occupational Choice over the whole state space (a, z)	156
3.13	Occupational Choice at increasing levels of the wage w	157

List of Tables

1.1	Summary statistics of demographic variables	20
1.2	Summary statistics of agricultural and administrative variables	21
1.3	Self-reported Shocks	22
1.4	Coping strategies	23
1.5	Food consumption (2016 \$) by categories	25
1.6	Different agro-ecological zones in the sample	26
1.7	SPEI values: categories	29
1.8	Descriptives various continuous indicators	30
1.9	Six months Spei distribution in December	31
1.10	Shocks correlations (across waves)	35
1.11	Effect on consumption of weather shocks: different continuous indicators	43
1.12	Coefficients of self reported shocks in the main specification (see Table 1.11)	44
1.13	Effect on consumption of weather shocks: different categorical indicators	45
1.14	Effect on consumption of weather shocks: categorical Spei in different months	45
1.15	Different coping strategies, continuous shock indicators	49
1.16	Different coping strategies, categorical shock indicators	50
1.17	Effect on consumption of weather shocks: continuous Spei in different months	52
1.18	Coefficients of the model estimated by grouping self-reported shocks into two categories.	52
1.19	Heterogeneity analysis: farm income terciles, agro-zones, long-run precipitations, and administrative regions	53
1.20	Effect on consumption of weather shocks: continuous indicators and spatial correlation	54
1.21	Effect on consumption of weather shocks: continuous indicators at a higher aggregation	54
1.22	Effect on consumption of weather shocks: continuous indicators and lags	55
1.23	Effect on consumption of weather shocks: levels	55
1.24	Effect on consumption of weather shocks: levels	56
1.25	Effect on consumption of weather shocks: levels	56
1.26	Different coping strategies, continuous shock indicators	56
1.27	Different coping strategies, categorical shock indicators	57
1.28	Effect on consumption of weather shocks: first differences	57
2.1	Industry-level summary statistics of OTT measures for goods industries only	82
2.2	Industry-level summary statistics of OTT measures for all industries	82
2.3	Transition matrix	106
2.4	SBS NACE Rev.1.1 – Raw Data	109
2.5	Number of missing observations within country-sector pairs	110
2.6	SBS Nace Rev.2 – Raw Data	110

2.7	Number of missing observations within country-sector pairs	111
2.8	Imputation on pre-2008 SBS	112
2.9	Imputation on post-2008 SBS sample	112
2.10	Production and concentration sample	122
3.1	Parameters of the static model	142
3.2	Parameters of the dynamic model	148

Overview

This thesis investigates a broad range of topics related to insurance, market power, and inequality, both from an empirical and a theoretical perspective. Apparently, the contents of this thesis are very heterogeneous: the first chapter is about consumption insurance in developing countries (Ethiopia in particular), the second is about industry concentration across selected OECD countries, and the third one is a note about a possible relationship between technological change and income inequality. As such, it may appear difficult to see a common motive; it turns out that heterogeneity is not only the final outcome (the contents of the thesis) but also what inspired it all.

In fact, each chapter is animated by the idea that heterogeneity in economics is pervasive but that there are situations in which looking at the aggregate behaviour of economic agents, when possible, is interesting (and even useful).

In the first chapter, I exploit the significant heterogeneity of the shocks hitting Ethiopian households and their heterogeneous response, using relatively recent data (World Bank's LSMS-ISA for households and satellite data for weather shocks). On the one hand, households seem able to insure against most idiosyncratic and mild adverse weather shocks. On the other hand, vulnerability to stronger weather shocks (especially droughts) remains elevated.

In the second chapter, starting from firms' individual data, aggregate trends about industry concentration and other proxies of competition are built.¹ The project innovates on the existing literature in its measurement of concentration, aimed at reflecting markets more accurately. On average, aggregate concentration is found to be increasing.

In the third chapter, which only lays out some preliminary steps of a more extensive inquiry, I model the heterogeneous effects of aggregate technological progress on individual economic agents and show how this can affect aggregate inequality and other aggregate indicators studied in the macroeconomics literature, such as the entrepreneurship rate and the overall firm distribution. It should be noted, however, that this note is a simple exposition of a possible modelling device rather than a full explanation of these phenomena.

¹This chapter is part of a larger project conducted at the OECD in the Productivity Innovation and Entrepreneurship Division of the STI Directorate, jointly with Christian Abele, Sara Calligaris, Chiara Criscuolo, Josh De Lyon, Andrea Greppi, and Miguel Chaves.

Chapter 1

Weather shocks and coping strategies in Ethiopia: a new indicator

1.1 Introduction

In this first Section, I introduce the topics touched on in this work. In the first part (Section 1.1.1), I give a rapid overview of the contents of this chapter, which also shows what motivated me in pursuing this subject; in the second one (Section 1.1.2), I review the literature that is closely related with this study and finally, in the third part (Section 1.1.3), I expose what are the main contributions of this work to the field.

1.1.1 Overview

Households in developing countries face a substantial amount of risk, which affects their welfare and their vulnerability to poverty (see Alderman and Paxson (1994), Morduch (1995), and Townsend (1995)). While households exhibit a certain ability to insure against shocks, previous studies showed that they are far from reaching the benchmark of perfect insurance, both from the intertemporal viewpoint and in terms of within-community risk sharing. For this reason, it remains highly relevant to study both aspects: the sources of risk on the one hand and the effectiveness of household strategies in dealing with shocks on the other. In fact, repeated exposure to adverse shocks can cause severe reductions in consumption and the inability of households to escape poverty (Carter and Lybbert (2012)).

Although the sources of risks can be manifold, one of the main threats comes from weather risks (Dercon (2004), Porter (2012)), particularly in environments where agriculture is the dominant source of income. The relative importance of weather shock for rural economies is one of many reasons for focusing on them. In the context of climate change, shocks can become more frequent and intense, posing threats to the growth of developing countries. Moreover, the availability of satellite data (Dell, Jones, and Olken (2014)) and new indicators for weather events might allow for more timely monitoring and precise assessment of weather shocks and, therefore, a higher ability in ex-ante prevention policy and better targeting of ex-post interventions.

To address these issues, at least locally, I focus on rural households in Ethiopia. The choice of Ethiopia is quite natural: the country is heavily dependent on agriculture, which is mostly rain-fed, and it is historically prone to weather fluctuations (see S. M. Hsiang and Meng (2015) for a more general assessment of tropical economies). For these reasons, Ethiopia is particularly vulnerable to weather shocks, and despite the high growth registered in the last decades, poverty rates remain high, especially in rural areas. While many studies focused on Ethiopia in the past, I differentiate from them by using new and more modern data, both at the household level (LSMS-ISA and CSAoE (2012)) and at the weather data level (Peng et al. (2020) and Funk et al. (2015)). In fact, one of the advantages of these new data is the possibility of matching information on both sides, using household geographical locations to obtain the history of weather shocks that they have undergone during the survey period.

By doing so, this work touches on two macro-areas of research in economics. The first one is about risk and coping strategies in developing countries and, from a theoretical viewpoint, on consumption smoothing across time and states of nature; the second one is a relatively new strand of research called climate economics (Dell, Jones, and Olken (2014), Auffhammer et al. (2013), S. Hsiang (2016)), which systematically tests the effects of weather and climate on economic outcomes. The main feature on which this literature relies is the fact that weather events can be considered, most of the time, exogenous to the economic outcomes that they influence.

I contribute to the literature by merging these two branches of research and applying climate economics tools (the use of satellite data) on household panel data to test consumption insurance theories and the ability of rural households in Ethiopia to deal with risk. In a certain sense, one of the contributions consists in providing new answers to relatively old questions: on the one hand, the test of consumption theories gave rise to an extensive literature that is still relevant nowadays. On the other hand, the possibility of relying on new and more precise data improves the quality of the answers.

This work provides two types of evidence of the ability of rural households to smooth consumption against shocks. On the one hand, households seem able to insure against most idiosyncratic and mild adverse weather shocks. On the other hand, vulnerability to stronger weather shocks (especially droughts) remains elevated. Further evidence shows that specific coping strategies (the possibility to smooth income through occupations different from the main agricultural one and the possibility to borrow from formal institutions) can alleviate the effects even of strong weather shocks.

A secondary result obtained through this work is the proof of the suitability of satellite data and in particular of data regarding the Standardized Evapotranspiration Index (SPEI) index (S. M. Vicente-Serrano, Beguería, and López-Moreno (2010), Peng et al. (2020)), to test consumption insurance theories and, more generally, to track household welfare in rural Ethiopia. Incidentally, this work is also one of the firsts using the SPEI index in economics (to the best of my knowledge, there are only other three: Azzarri and Signorelli (2020), Albert, Bustos, and Ponticelli (2021), and Piolatto et al. (2022)). This ability to track weather conditions timely can be very relevant in a shock-prone country such as Ethiopia, especially under the additional risk of climate change (see United Nations (2012)).

1.1.2 Literature Review

In this Section, I review the literature, focusing on the contributions that are the closest to the present work. I divided the Section into three parts: the first one discusses works relative to the theory of consumption insurance, and by mentioning the works that laid the conceptual foundations

of the field, it has a more historical flavour; the second one, by talking about applications and testing of the theories exposed in the previous one, especially in the context of developing economies, is more closely related with the present study; the third one exposes some of the works from the field of climate economics.

Consumption smoothing

The literature on consumption smoothing is very extensive. Therefore I only provide a selection of references that have a historical relevance for the field or that are more closely related to the present work. For a more detailed description of this field, please refer to the literature reviews by Jappelli and Pistaferri (2010) (who have also written a more extended textbook version, see Jappelli and Pistaferri (2017); for another textbook treatment, refer to Deaton (1992)), by Attanasio (1999), by Attanasio and Weber (2010), and by Meghir (2004).

The two cornerstones of this literature can be considered the studies, respectively, by Modigliani and Brumberg (1954) and by Friedman (1957). The first one exposed the well-known life-cycle model. In contrast, the second one proposed the permanent income model as an alternative to the, back then, traditional Keynesian consumption function¹. In both works, the idea is that consumption should react only slightly to anticipated changes in income and that economic agents would use their savings to smooth consumption.

More generally, the whole theory of consumption smoothing can be thought of as an explanation of what happens to consumption when the amount of available resources is subject to a change. Following the broad distinction outlined in Jappelli and Pistaferri (2010), the literature evolved into two separate (although very close) branches: the first one investigated the effect on consumption of anticipated changes of income, while the second, which is closer to the present work, studied the effects of unanticipated shocks².

One of the first important results of the former line of research was obtained by Hall (1978), who showed how, under certain assumptions, consumption (or, more generally, marginal utility) follows a martingale process: this means that anticipated variation in income should not affect consumption when they occur, since agents expectations, and therefore consumption levels, already incorporate that information. Hall (1978) results were further tested by Flavin (1981) and by Campbell (1987), who found what was called “excess sensitivity” of consumption to anticipated income growth. The literature evolved by chasing explanations for these failures of the theory (binding liquidity constraints, preferences specifications such as leisure consumption non-separabilities or habit persistence, and the presence of durable goods; please refer to Jappelli and Pistaferri (2017) for a detailed description of each of these features), and by differentiating between positive anticipated changes in income and negatives one. However, these advancements are too specific to be discussed here.

The second macro-area in the field of consumption smoothing, which is more relevant to the present study, investigates the effects on consumption of unanticipated changes in income. According to Jappelli and Pistaferri (2010), the literature followed three distinct approaches in pursuing this line of investigation. The first one, which Jappelli and Pistaferri (2010) calls the “quasi-experimental approach”, seeks to identify situations where income experiences unforeseen changes and assess how consumption responds to these shocks within a quasi-experimental setting. The

¹In fact, the work contained in Ramsey (1928) should be probably considered as the real ancestor of the modern theories of consumption and savings, as nicely pointed out in Attanasio (2015).

²The other distinction, somehow posited already by Friedman (1957), is between permanent and transitory shocks.

second evaluates the marginal propensity to consume in response to income shocks by imposing the consumption-income covariance restrictions that emerge from theory (see Hall and Mishkin (1982) for an early contribution using this method). The third one exploits the differences between agents' expectations and actual outcomes to identify income shocks (see Manski (2004) for a theoretical explanation).

In this literature review, I focus only on the first approach, the quasi-experimental one, since it is the one applied in this work. I provide a simple theoretical exposition of this approach in Section 1.3, where the conceptual framework used in this study is presented. This method does not require estimating an income process; instead, it compares households affected by shocks with those unaffected (or the same household before and after the shock occurs). Jappelli and Pistaferri (2010) mention the paper by Bodkin (1959) as an example of pioneering work where all the elements of this approach are present: consumption reaction is tested after an unexpected dividend of National Life Insurance is obtained by World War II veterans. An early use of such an approach that used climate variables can be found in Wolpin (1982) and Paxson (1993), who use weather shocks to identify the exogenous income variation and its effect on consumption. The next Section describes more examples of works that adopted such an approach for studying consumption insurance in developing economies.

risk in developing countries

Since this paper is about consumption insurance of Ethiopian rural households, this Section is devoted to a description of the literature that explored the impact of shocks on consumption in rural economies. The literature can be divided into three large strands. The first one is devoted to tests of full insurance and traces back to Townsend (1994), who explored the extent to which rural households in India could insure their consumption. The theoretical benchmark is a complete market economy: such an environment represents a village economy, where informal arrangements are made and have the purpose of insuring households of the village against idiosyncratic shocks. The empirical counterpart, which tests whether household consumption variations are orthogonal to household income variations, is the so-called complete market test, or Townsend test. The paper by Townsend (1994) contained a result that was considered surprising at that time: although full insurance was rejected, "Household consumptions comove with village average consumption. More clearly, household consumptions are not much influenced by contemporaneous own income, sickness, unemployment, or other idiosyncratic shocks, controlling for village consumption (i.e. for village level risk)." A whole body of research grew out of Townsend's results, both theoretical and empirical: the former investigated mechanism that could explain the failure of insurance (presence of moral hazard and limited commitment; see Kocherlakota (1996)), the latter tested in the data the restrictions produced by the different models explaining the lack of insurance (Ligon (1998), Kinnan (2022)).

The second strand of literature tested numerous empirical formulations of the Permanent Income Hypothesis (Friedman (1957), Meghir (2004)), studying whether and how rural households are able to intertemporally insure their consumption when shocks hit them. For general reviews of the literature, please refer to Alderman and Paxson (1994), Morduch (1995), and, for a textbook exposition, consider Bardhan and Udry (1999). Early studies such as those by Paxson (1992), Wolpin (1982), and Rosenzweig and Wolpin (1993) showed how rural households used asset dis-saving strategies to insure against transitory shocks. A different result was obtained by Kazianga and Udry (2006), who found that usual coping strategies such as risk sharing and the use of assets as buffer stock were not effective and that the amount of consumption smoothing was very low in the context of the 1981-1985 drought in Burkina Faso.

A more direct approach was adopted by Dercon (2004), by Dercon, Hoddinott, and Woldehanna (2005), and by Porter (2012), who tested the effects of poor rainfall and other shocks on consumption growth in Ethiopia. While they all found that negative weather shock affected consumption, idiosyncratic shocks did not have an impact on it in the study by Porter (2012). This result suggests a partial ability of households to insure themselves. Porter (2012) also found that only harder weather shocks affected consumption, while households were able to insure against milder shocks. This fact is in line with the results of the present work, which suggest both the ability to insure against mild weather shocks and the presence of effective coping strategies, such as income smoothing through the diversification of labour supply (ex-ante coping strategy) and the use of formal credit markets (ex-post coping strategy).

Two more recent works, by Gao and Mills (2018) and by Colmer (2021), show how the literature in the field of consumption insurance is well-alive. Both studies focus on rural households in Ethiopia, and the former finds that off-farm employment and the presence of formal safety net transfers can partially dampen the effects of adverse shocks on household consumption. In contrast, migration, remittances, and informal safety nets are not effective. The latter focuses on household well-being, finding that it is affected by rainfall variability. Also, these results are partially in line with the present work, where both informal and formal safety nets are shown to be ineffective in softening the impact of adverse weather shocks. Other contributions adopting such a direct approach are those by Skoufias and Quisumbing (2005), who report results for different countries, showing how the reduction in non-food consumption is sometimes used to smooth food consumption. Another testimony of the multitude of studies conducted using data from studies about Ethiopia can be found in Demeke, Keil, and Zeller (2011), who tests the effect of weather shocks on food security and shows how rainfall is a crucial determinant of it. Finally, Hill and Porter (2017) study the vulnerability of Ethiopian households to droughts and increases in food prices and find that many households are unable to protect against such shocks.

Climate economics

The present work can be linked with a relatively recent strand of literature called climate economics (Dell, Jones, and Olken (2014)). This literature exploits econometric techniques, often panel data, to evaluate how the influence of various weather and climatic characteristics (precipitations, temperatures, storms) affect economic and social outcomes. These methods leverage the fact that weather variables can be used as a source of exogenous variation, and therefore, the identification strategy of the effects is usually quite clean. The link between the present study and this literature is quite clear: I exploit weather data to study rural household welfare in Ethiopia and their ability to cope with climatic risk.

The literature expanded greatly in the last decade, and it is useful to refer to several articles that systematised its findings. The review by Dell, Jones, and Olken (2014) contains a quite general overview of the field, which presents its most common research designs and provides an overview of the types of data typically used (this is also provided in Auffhammer et al. (2013), who also discuss potential pitfalls in using them). It is important to mention that there are several types of weather data: ground station data, gridded data, and satellite data³. Ground station data refers to the recording of various variables by isolated weather stations. Gridded data interpolates such recordings to obtain a grid of observations of the relevant variables. Satellite data use satellite monitoring to extrapolate information about the variables. The present work uses CHIRPS Data for precipitations (see Funk et al. (2015)) and SPEI data (see S. M. Vicente-Serrano, Beguería,

³Dell, Jones, and Olken (2014) also mention reanalysis data, that relies on climate models that use combinations of the various data types in order to obtain estimates of the needed weather variable.

and López-Moreno (2010) and Peng et al. (2020)), which are obtained through a combination of the former methods. I refer to S. Hsiang (2016) for other technical subtleties, which contain several valuable details. There, the difference between climate and weather is exposed, and several challenges and solutions to issues arising in climate econometrics.

In terms of economic contributions, according to Dell, Jones, and Olken (2014) and Carleton and S. M. Hsiang (2016), weather data have been used to study very different topics and, in particular, the impact of weather on agricultural yields (Deschênes and Greenstone (2007), labour productivity, trade, energy supply and demand, health, conflicts, and migrations. Another direction followed the study of the impact of natural disasters on economic outcomes.

Finally, a few other examples of works that are close to the present study (and might have ended in the previous Section of the literature review, but since they explicitly use weather data, I decided to mention them here). Two papers using subsets of the Ethiopian data that I am using (asking different questions, though) are by Auci, Castellucci, and Coromaldi (2018) and Auci and Coromaldi (2021). The other is by Azzarri and Signorelli (2020), who study the spatial determinants of welfare in Sub-Saharan Africa using the SPEI index to account for climatic variability. This is the only work, with those by Albert, Bustos, and Ponticelli (2021) and Piolatto et al. (2022), using the SPEI index in economic literature.

1.1.3 Contributions

The contributions of this paper can be divided into two categories: a technical one and another with more relevance to economics.

From a technical viewpoint, I create a new dataset by merging two fairly new and independent data sources. The first piece of data is the survey panel LSMS-ISA for Ethiopia over the period 2011-2016, provided by the World Bank and the Central Statistical Agency of Ethiopia (see Section 1.2.1 for further explanation), which contains highly detailed information on rural households in Ethiopia (in particular, and most crucially, their geolocation). The second piece of data is given by two satellite datasets on weather indicators: precipitations and the Standardized Evapotranspiration Index (SPEI) index (S. M. Vicente-Serrano, Beguería, and López-Moreno (2010), Peng et al. (2020)).

Adopting the SPEI index as a weather indicator represents already a minor contribution since this is one of the first papers in economics that exploits it (the others are, to the best of my knowledge, Azzarri and Signorelli (2020), Albert, Bustos, and Ponticelli (2021), and Piolatto et al. (2022)), and the first that uses it in testing consumption insurance theories.

More importantly, building this dataset allowed me to have more detailed and more suited data for my purposes: previous works on risk and insurance in Ethiopia (for example, Dercon (2004), Demeke, Keil, and Zeller (2011), Porter (2012), Gao and Mills (2018), and Colmer (2021)), relied on data gathered in a small number of villages across the country. Although these villages had been selected to guarantee a minimum of representativeness in terms of geographical characteristics, the consequence was a substantial reduction in between variation across the observations. Although this, in principle, might be innocuous when studying the effects of aggregate shocks, the (too) small amount of variation might hide interesting effects.

From the economics category viewpoint, this work provides two types of evidence of the ability of rural households to smooth consumption against shocks. On the one hand, households seem able to insure against most idiosyncratic and mild adverse weather shocks. On the other hand, vulnerability to stronger weather shocks (especially droughts) remains elevated. Further evidence

shows that specific coping strategies (the possibility to smooth income through occupations different from the main agricultural one and the possibility to borrow from formal institutions) can alleviate the effects even of strong weather shocks.

1.2 Data

This Section describes the socioeconomic context of this study, the variables used, and the sources from which they are gathered. The Section is organised in three parts. In Section 1.2.1, I describe the information retrieved from the Ethiopian Rural Socioeconomic Survey (ERSS). First, I focus on household characteristics, which include demographic information at the household level, several indicators about agricultural practices, information on shocks, coping strategies, and household geographical locations. Then, I describe how the measure of household food consumption is defined and built. This variable is crucial for this study since it is the proxy to capture household welfare. In Section 1.2.2, I describe the data sources of the weather indicator used in this study: cumulative precipitations over a season and the Standardised Precipitation-Evapotranspiration Index (SPEI). Finally, in Section 1.2.3, I describe the procedure adopted to match the two different data sources: I exploit the information on household geographical locations in order to have, for each household, measures of weather shocks that occurred across the three waves of the survey.

1.2.1 Households Survey

Households data are gathered by means of the Ethiopian Rural Socioeconomic Survey (ERSS) for the first wave (2011/2012, please refer to LSMS-ISA and CSAoE (2012)) and by means of the Ethiopian Socioeconomic Survey (ESS) for the second and third wave (2013/2014, please refer to LSMS-ISA and CSAoE (2015) and 2015/2016, see LSMS-ISA and CSAoE (2017)). The various surveys are the product of the combined effort of the Central Statistical Agency of Ethiopia (CSA) and the World Bank Program Living Standard Measurement Study - Integrated Surveys of Agriculture (LSMS-ISA). LSMS program was born with the aim of studying living standards and inequality through representative household surveys. ISA is a project that targets countries among the poorest in the world and is particularly well suited for measuring agricultural income, which is the primary source of income in such countries⁴.

In the first wave, the survey was designed to be representative of rural and small town areas of Ethiopia (as reported in LSMS-ISA and CSAoE (2012)), according to a two-stage probability sample⁵. In the first stage, 290 enumeration areas (EAs) were selected to represent rural areas and 43 EAs to represent small towns. In the second stage, about 12 households were selected in each EA. A total of 3969 households out of the 3996 originally selected were interviewed, with a response rate of 99.3% (out of these households, 3108 became part of the final sample of the first wave for this study due to data quality issues).

In the second and third wave, the structure of the survey remained substantially unaltered except for the addition of 1500 urban households to the sample⁶. The follow-up rate (LSMS-ISA and CSAoE (2015)) of the households interviewed for the first wave is 95% (3776 households, out of which 3137 are part of the final sample of this study for the second wave). Finally, in the third wave (LSMS-ISA and CSAoE (2017)), the total number of households interviewed in the first wave with complete interviews is 3699 (out of which 3248 are part of the final sample of this study).

Notably, households were geolocated. As per the survey documentation (LSMS-ISA and CSAoE (2012)), in order to grant data confidentiality, "for small towns and urban areas, an offset range of 0-2 km is used. In rural areas, where communities are more dispersed and the risk of disclosure

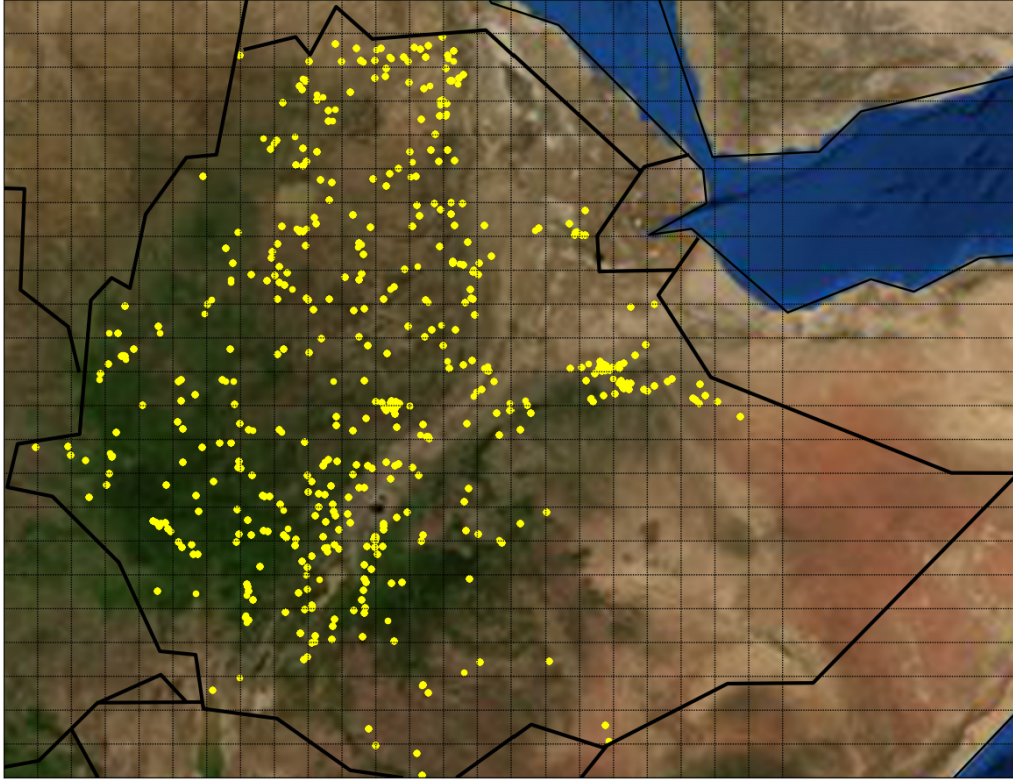
⁴Countries currently with LSMS-ISA datasets: Burkina Faso, Ethiopia, Malawi, Mali, Niger, Nigeria, Tanzania, Uganda.

⁵According to the CSA any town below 10000 habitants is considered to be small.

⁶In this study I focus exclusively on rural households, and therefore I exclude from the final sample the urban households added from the second wave on.

may be higher, a range of 0-5 km offset is used. Additionally, an offset range of 0-10 km is applied to 1% of EAs”. For the purposes of this study, such a randomisation procedure does not pose any limit to the analysis. In Figure 1.1, the location of the households is plotted. As it can be seen, and in a crucially different way than previous studies, households are scattered throughout the whole territory of Ethiopia (apart from the desertic Somali region in the southeast of the country).

Figure 1.1: Survey Enumeration Areas distribution over Ethiopia




The survey comprised five questionnaires (please refer to LSMS-ISA and CSAoE (2012) or other waves’ documentation for further details).

- A *household questionnaire*: it was administered to all household members. It collected information on demographics, use of time and labour, savings, expenditures on food and non-food items, information on non-agricultural income-generating activities, shocks and food security, safety nets, housing conditions, assets, credit, and other sources of income.
- A *community questionnaire*: it was administered to a group of community members to collect information on the socio-economic conditions of the enumeration areas where the sample households reside and how they perceived such conditions evolved
- Three *agriculture questionnaires*: they consisted of a post-planting (PP) questionnaire, post-harvest (PH) questionnaire and livestock (L) questionnaire that were administered to all household members who are engaged in agriculture activities

The post-planting questionnaire was administered between September and October, the livestock

Figure 1.2: Survey timing across waves



questionnaire took place between November and December, and the household, the community, and the post-harvest questionnaires took place between January and March of the following year⁷. Figure 1.2 is a synthetic representation of the survey administration and agricultural-meteorological timing across surveys.

Household data

While some information gathered through the surveys is available at the individual level, the unit of analysis in this study is the household. In this work, I focused exclusively on rural and small-town households. Due to data limitations (either for consumption data or for weather data), some households are excluded from the final sample, which consists of 9493 observations across the surveys, with 3108 households retained in the first wave, 3137 in the second wave, and 3248 in the third one. I characterise households through many variables: some refer to household composition, some to characteristics of the household head, some others capture agricultural practices of the family, geographical information, and exposure to shocks.

Table 1.1 reports descriptive statistics about the demographic characteristics of the households, averaged across waves. Households have an average of 5.7 members. About 24% of the families have a female head, and the average age of the head is 45 years. The average family age is about 23 years, and the dependency ratio (defined as the number of non-working members over the number of working members) is about 1.3. These numbers are reassuring for the quality of the data since they are very similar in magnitude to those of other studies (that used other data from Ethiopia) such as those by Demeke, Keil, and Zeller (2011), Gao and Mills (2018), and Porter (2012).

Table 1.1: Summary statistics of demographic variables

Variable	Mean	Std. Dev.	Min.	Max.	N
Dep. Ratio (non working / working)	1.34	1.09	0	11	9493
Household head is female	0.24	0.43	0	1	9493
Household's head age	45.13	14.55	13	100	9493
Household's head can write and read	0.42	0.49	0	1	9493
Household Size	5.69	2.48	1	17	9493
Household mean age	23.66	10.27	6.75	80	9493

⁷The first wave of the Ethiopian survey took place between 2011 and 2012: for this waves the months between September and December refer to 2011, the months between January and March refer to 2012. Please notice that for the third wave, some minor modifications in timing took place; see LSMS-ISA and CSAoE (2017) for further details.

In Table 1.2, I report the descriptive statistics of the variables to which I refer to as agricultural-geographical variables. In the study sample, 89% of the households conduct agricultural activities; they have 2.4 Tropical Livestock Units (TLU) and 0.5 hectares of cultivated land on average. These two variables are crucial since they serve as proxies for permanent income in the econometric analysis (the use of TLU as permanent income is relatively standard in studies of consumption insurance in developing countries, see, for example, Demeke, Keil, and Zeller (2011), Gao and Mills (2018), and Porter (2012)). In terms of technologies and coping strategies adoption, on average, 75% of the households cultivate more than one crop type (4.69 different species on average), about 40% have some vaccinated livestock units or use inorganic fertilisers, but only 15% use any improved seeds. The last four variables in Table 1.2 show that household locations are relatively remote with respect to administrative or economic points of interest.

Table 1.2: Summary statistics of agricultural and administrative variables

Variable	Mean	Std. Dev.	Min.	Max.	N
Any agricultural activity	0.89	0.31	0	1	9493
TLU survey year	2.42	3.38	0	69.3	9493
TLU 1 year before	2.75	3.87	0	71.3	9493
Any vaccinated animals	0.37	0.48	0	1	9493
Multiple crops	0.75	0.43	0	1	9493
Effective number of crop species	4.69	3.41	0	37.1	9493
Any improved seeds	0.15	0.36	0	1	9493
Any inorganic fertilizer	0.39	0.49	0	1	9493
Cultivated area (Ha)	0.95	5.60	0	426.51	9493
Agriculture within approx 1 km (%)	32.34	20.23	0	97	9493
Plot Distance (Km)	1.22	13.88	0	775.7	7994
Annual Mean Temperature (°C)	19.32	3.44	10.2	29.4	9493
Annual Precipitation (mm)	1097.45	395.55	144	2031	9493
Precipitation of Wettest Month (mm)	226.63	78.88	39	448	9493
Precipitation of Wettest Quarter (mm)	569.03	213.48	80	1184	9493
Mean Temp. of Wettest Quarter (°C)	19.07	3.66	10.3	31.9	9493
Avg. rainfall for Jan-Dec (mm)	881.69	287.28	247	1696	9493
Market distance (Km)	68.12	51	0.3	283.3	9493
Adm. center distance (Km)	169.47	129.39	1	773.1	9493
Nearest City (pop.>20,000)	38.85	30.11	0	208.2	9493
Nearest Border Crossing (Km)	252.67	108.76	8.5	501	9493

Self-reported shocks and coping strategies

Table 1.3 reports the average occurrence of shocks across waves, self-reported by the households, as they were gathered in Section 8 of the *Household questionnaire* (see LSMS-ISA and CSAoE (2012)). The most reported ones are price rise of food items (20%), drought (19%), illness of a household member (16%), and increase in the price of inputs (11%). As a robustness check, I also built an indicator for generic idiosyncratic and aggregate shocks⁸. Please notice that while climate shock indicators are present as self-reported shocks, I decided not to use them in the present work

⁸I gathered the death of the bread-winner, illness of household member, loss of non-farm job, theft/robbery/violence, loss of house/farm/land, displacement (govt. project), death of other household member into the idiosyncratic shocks indicator, and drought, flood, landslides/avalanches, heavy rain preventing work, price fall of food items, price rise of food items, increase in the price of inputs.

and rely exclusively on the objective data obtained from satellites.

Table 1.3: Self-reported Shocks

Variable	Mean	Std. Dev.	N
death bread winner-head	0.02	0.15	9493
illness Household member	0.16	0.36	9493
loss non farm job Household member	0.01	0.1	9493
drought	0.19	0.39	9493
flood	0.02	0.15	9493
landslides/avalanches	0	0.07	9493
heavy rain preventing work	0.02	0.15	9493
other crop damage	0.06	0.24	9493
price fall of food item	0.03	0.17	9493
price raise of food item	0.2	0.4	9493
increase in price of inputs	0.11	0.31	9493
great loss/death livestock	0.07	0.25	9493
fire	0	0.07	9493
theft/robbery/violence	0.01	0.09	9493
loss of house/farm/land	0	0.07	9493
displacement (govt. proj)	0	0.03	9493
local unrest/violence	0.01	0.09	9493
other	0.02	0.13	9493
death other Household member	0.02	0.13	9493
idiosyncratic shock	0.21	0.41	9493
aggregate shock	0.38	0.49	9493

While the surveys report the coping strategies adopted by households, I also build indicators that report whether families adopted some mechanisms that, in principle, might have softened the effect of shocks. These indicators are not entirely mutually exclusive, so their effectiveness in dampening the shocks will be tested once at a time in Section 1.4.3⁹. In Table 1.4, I report the average values across waves of such indicators.

The first indicator shows that 18% of the observations received a transfer from a Government or Non-Governmental Agency or Program. I refer to such indicator as the presence of “any formal safety net”. Such a variable is a rather crude indicator since it gathers the presence of any transfer from such institutions, as recorded in Section 13 (Assistance) of the *Household questionnaire*. These transfers can be of different natures, and the categories used in the questionnaire are: i) free food, ii) Food-for-work programme or cash-for-work programme, iii) Inputs-for work programme, iv) other. A separate category is any transfer from the Productive Safety Net Program (PSNP) of Ethiopia. PSNP is a national-scale program that provides a safety net for chronically insecure and poor households. Transfers can be conditional on work participation or unconditional. In section 13 of the questionnaire, only unconditional transfers from PSNP are registered. Payments against public work are registered in Section 4 (and result anyway as participation to PSNP, see the third indicator in Table 1.4).

In contrast, the variable “any informal safety net” captures the fact that about 16% received a transfer from friends or relatives. It is built using information gathered through Section 12

⁹For example, the presence of any formal safety net partly overlaps with the participation of the households to the PSNP program.

(Other income) of the *Household questionnaire*. Households are asked whether they received any gift or transfer (cash, food, or in-kind) and to attach a monetary value to them. I decided to use an indicator to signal the presence of so-called informal safety nets. The presence of such mechanisms has been widely studied and debated, and they are considered one of the strategies that households in developing countries have to insure themselves against shocks (see, for example, Townsend (1995)).

The third indicator accounts exclusively for the households' participation in the PSNP program described above. In this sense, it includes not only unconditional transfers granted by the program but also other initiatives such as food-for-work programs.

The variable, "any off-farm work hours", indicates that 12% of the observations (households across waves) had members who complemented their agricultural income by shifting their labour to off-farm types of occupations; from an economic viewpoint this indicates the ability of income smoothing in order to diversify risk (see Kochar (1999), Morduch (1995), or Porter (2012) as references)

The variable "any ex-ante (agric.)" shows that the majority of observations with agricultural activity adopted at least one of the following technologies: the use of improved seeds, the use of fertilisers, or the cultivation of several types of cropping. I refer to these practices as ex-ante coping strategies because, in principle, they should help rural households become more resilient to shocks. An interesting contribution about the ability of Ethiopian households to adopt such technologies is the one by Dercon and Christiaensen (2011).

The last coping strategy adoption indicator refers to the use of formal credit instruments, in particular the use of mortgages and borrowings from banks, and it is a piece of information gathered in Section 4 and Section 11b of the *Household questionnaire*. Although this refers also to the sources of credit used in non-farm businesses for some households, it may be a general indication of access to financial services. About 12% of the observation has access to some form of borrowing. This can be regarded as an ex-post coping strategy, since in the presence of well-functioning credit markets, households borrow in order to smooth consumption.

Table 1.4: Coping strategies

Variable	Mean	Std. Dev.	N
any formal safety net	0.18	0.39	9493
any informal safety net	0.16	0.37	9493
participation to PSNP	0.11	0.32	9488
any off farm work hours	0.12	0.32	9493
any ex-ante (agric.)	0.76	0.43	9493
formal borrowing	0.12	0.32	9493

Food consumption data

The variable chosen to capture household welfare is food consumption. To build the food consumption aggregate, i.e. the total annual expenditures of a household on food, I mostly follow the procedure described by Magalhães and Santaella-Llopis (2018). Also, the original dataset provides a consumption aggregate, and I exploit some of the procedures described in the relative documentation to build the measure adopted in this work. The issue with the consumption aggregates originally provided with the raw dataset is that they are not comparable across waves. Every wave's aggregate is built using the prices of that period. Therefore, any variation of consumption between two waves, using such aggregates, could be attributed either to variation in quantity or

to variation in prices. To meaningfully compare consumption across waves, common prices (e.g., only those of a given wave) must be used to build consumption aggregates. For this reason, the procedure used to build consumption aggregates and to extract prices of the food items composing the consumption basket is crucial.

Section 5A of the survey asks households about food consumption in the previous seven days. For each food item of a given list¹⁰ they are asked about the consumed quantity and its sources: whether it came from purchases, own production, or gifts). Each item comes with its unit of measure, which is possible to convert into kilograms with the aid of conversion tables provided with the data. Since households are asked how much they spent for the purchased amount of every item, it is possible to retrieve the prices per kilogram¹¹.

Using prices per kilo, it is possible to assign a monetary value also to quantities coming from households' own production or in-kind transfers and gifts. Since for certain households some food items came exclusively from their own production or in-kind transfers and gifts, it is not always possible to apply directly the procedure that I described in the previous paragraph. In those cases, I used the median price at the lowest possible geographical aggregation level for such food item, as long as at least ten observations are available. If they are not, I compute the median price of such a food item at a higher level of geographical aggregation, and so on¹².

Once all prices were retrieved, I computed the food consumption aggregate using the third wave's prices. In this way, for example, a kilogram of Teff consumption by a certain household registered in the first wave is evaluated at the third wave price. I also apply regional deflators provided in the original data to make expenditures comparable across regions. To obtain a yearly measure, I multiplied by 52 the aggregate obtained. Finally, prices were converted into 2016 US Dollars.

In Table 1.5, I report descriptive statistics of yearly food consumption at the household level along several dimensions, in particular those used to split the sample during robustness checks in the empirical section (see Section 1.4.2). Consumption is increasing in income terciles, and this is reassuring. The second and third categories, essentially climatic ones, show that regions with lower rain on average have higher consumption¹³. In terms of administrative subdivision, it can be seen that Harari, Dire Dawa, and Afar regions have higher average food consumption. Across waves, average food consumption increases between the first and the second waves and decreases between the second and third waves. Such a pattern suggests that the considerable drought that hit the country between 2015 and 2016 (with effects lasting even further in the following years) is captured in the data. Finally, the last row of Table 1.5 gives the average value of yearly food consumption at the household level in 2016 US Dollars across waves. Such an average value of about 910 (2016) US Dollars appears to be reasonable. A back-of-the-envelope calculation using World Bank aggregate data suggests that Ethiopia's average household consumption expenditure in 2016 was between 2500 and 3000 US Dollars. Since the sample of this study only uses rural households (whose consumption is lower on average than that of urban households) and focuses on food consumption and not on total consumption, the average value obtained for the period

¹⁰most representative items: Teff, Wheat, Barley, Maize, Sorghum, Millet, Horsebeans, Field Pea, Chick Pea, Lentils, Haricot Beans, Niger Seed, Linseed, Onion, Banana, Potato, Kocho, Milk, Meat, Cheese, Eggs, Sugar, Salt, Coffee, Chat / Kat

¹¹The quantity purchased of every item in each wave was trimmed at the 97th percentile (using all non-zero values as distribution reference).

¹²The different levels of geographical aggregation used to retrieve median prices, in increasing order of aggregation, is Enumeration Area, Kebele, Woreda, zone, and national.

¹³As it is explained in the next section, the arid non-arid subdivision is based on a climatic classification provided with the raw data (see 1.4 for a visual representation). The second one instead is based on long-run average precipitations.

2011-2016 appears to be reasonable¹⁴.

Table 1.5: Food consumption (2016 \$) by categories

Income tercile	Mean	Std. Dev.	N
first	730.74	622.52	2953
second	840.39	645.65	3235
third	1139.95	864.03	3305
Agro-zone			
non-arid	814.67	625.42	5646
arid	1051.33	869.52	3847
LR precip. tercile			
low rain	1129.63	927.7	3108
average rain	812.95	600.31	3274
high rain	794.47	612.91	3111
Administrative region			
Tigray	821.1	552.19	1003
Afar	1219.76	1102.95	218
Amhara	672.26	484.32	2014
Oromia	1055.27	699.34	1911
Somalie	1484.7	1322.71	577
Benshagul Gumuz	533.22	326.93	350
SNNP	788.74	628.78	2455
Gambelia	807.65	637.55	294
Harari	1638.05	972.17	334
Dire Dawa	1245.76	649.24	337
Wave			
first	958.15	845.09	3108
second	933	735.07	3137
third	843.39	634.53	3248
Overall			
food consumption	910.57	743.28	9493

1.2.2 Weather data

Ethiopia is a highly diverse country from a geographical perspective. In terms of latitude, it is located within the tropical zone (3–15 degrees North in terms of latitude), and its territory exhibits a great range of topographical and climatic variability. Such geographical diversity, which is also due to the altitude variation (elevation ranges from about 100 meters above the sea level to more than 4500), is portrayed in Figure 1.4, where the different Enumeration Areas are classified in terms of Agro-Ecological-Zones (data provided with the Ethiopian Rural Socioeconomic Survey)¹⁵. Table

¹⁴I obtained these number by dividing the final consumption expenditures in 2016 (54.46 billion Dollars) by the total population in 2016 (105293228), and later by multiplying that value by the average family size (5.1, although urban households are smaller in size)

¹⁵As reported in Sebastian (2009): "Agro-ecological zones (AEZs) are geographical areas exhibiting similar climatic conditions that determine their ability to support rainfed agriculture. At a regional scale, AEZs are influenced by latitude, elevation, and temperature, as well as seasonality and rainfall amounts and distribution during the growing season. The resulting AEZ classifications for Africa have three dimensions: major climate zone (tropics or subtropics), moisture zones (water availability) and highland/lowland (warm or cool based on elevation)".

1.6 reports the number of observations per agro-ecological zone. In the robustness section of the empirical results (Section 1.4.2), I group the observations into two categories (arid or non-arid) based on their agro-ecological zone belonging.

Table 1.6: Different agro-ecological zones in the sample

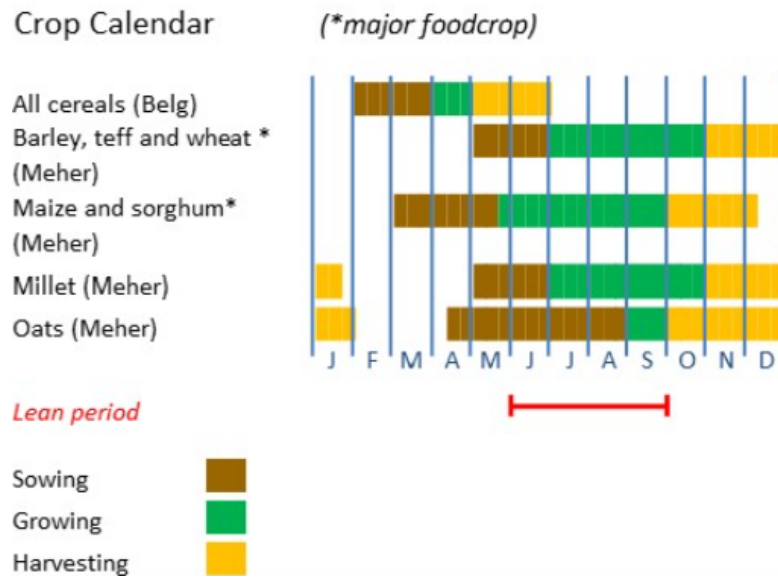
Agro-ecological Zones	Freq.	Percent	Cum.
Tropic-warm/arid	102	1.07	1.07
Tropic-warm/semiarid	620	6.53	7.61
Tropic-warm/subhumid	297	3.13	10.73
Tropic-warm/humid	14	0.15	10.88
Tropic-cool/arid	34	0.36	11.24
Tropic-cool/semiarid	3,091	32.56	43.80
Tropic-cool/subhumid	3,774	39.76	83.56
Tropic-cool/humid	1,561	16.44	100.00
Total	9,493	100.00	

Ethiopia is still largely a rural country that depends on agriculture (and where 95% of production is made by small-holder farmers; see Wakjira et al. (2021)), which accounts for more than 30% of its GDP and more than 70% of its exports (Negeri (2017a)). Agriculture, in turn, is primarily rain-fed and, therefore, is highly dependent on weather outcomes.

There are three climatological seasons in Ethiopia. Bega is the long dry season that goes from September to February. Belg is a short rainy season that takes place between February and April. Finally, the long rainy season, which is the most important from an agricultural viewpoint, takes place between June and early September, and it is called Kiremt¹⁶. In this study, I will focus primarily on shocks affecting the long agricultural season, which accounts for most of the agricultural production across the country (more than 90%, according to Ahmed, Tesfaye, and Gassmann (2023)). Figure 1.3 (obtained from FAO (2023)) illustrates the timing of major cropping in Ethiopia. Such a high dependence on Meher season is what led me to the adoption of the 6-month Standardized Evapotranspiration Index (SPEI) (which, as explained in Section 1.2.2, is a measure of cumulative rainfall deficits/surpluses particularly suited to track wetness conditions over a given time period) indicator to track weather conditions.

¹⁶The agricultural season associated with such rainy season is called Meher, and I will use these two terms interchangeably (as it is done in common parlance, as far as I understand).

Figure 1.3: Crop Calendar in Ethiopia



Source: FAO, 2023

Extreme weather events and general weather variability proved crucial to people's welfare (see Dercon (2004) for an early evaluation of the effect of rainfall variability on consumption growth in Ethiopia), and more specifically for food security. The ability of households to cope with weather risk becomes increasingly important given the recent trends in the occurrence of extreme events, which appear to be triggered by climate change.

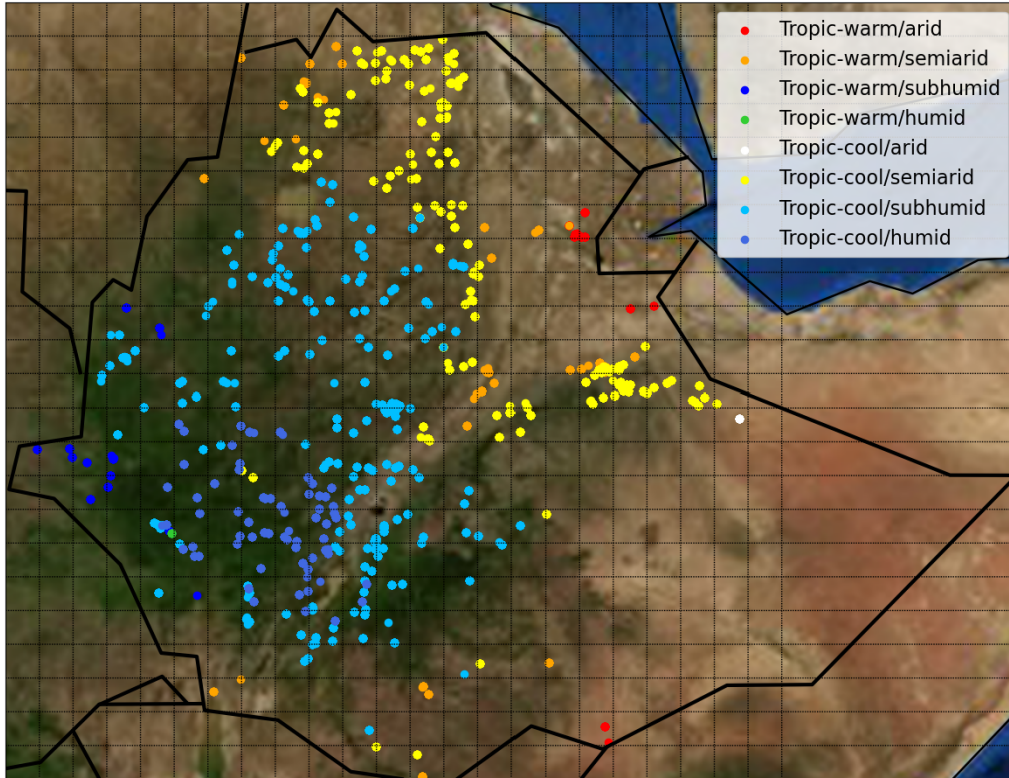
In particular, a relatively frequent event is drought. Droughts can be localised but also hit the majority of the national territory and can be exacerbated by El Niño cycles, such as the 2015-2016 drought (see Negeri (2017b))¹⁷.

Previous studies that tried to evaluate the effects of weather shocks on household consumption and welfare suffered from two limitations. First, they relied only on precipitation anomalies to indicate weather shocks. These are only sometimes sufficient indicators since pre-existing conditions of soil moisture with respect to a rainfall deficit can determine the onset of an agricultural drought. For this reason, I use a different indicator, the Standardized Evapotranspiration Index (SPEI). A second limitation is due to a combination of characteristics of household data and weather data: on the one hand, households were surveyed in a limited number of villages across Ethiopia, leading to a limited amount of between-variation across observations; on the other hand, climate data came from the weather station close to the village, leading to a further reduction in the available variation (with the exception of the study by Gao and Mills (2018)), and possibly, to measurement error in the actual weather conditions that affected the households. Using novel indicators (such as the SPEI) and satellite data that can be matched more closely to household locations should mitigate such limitations. In the following two subsections, I describe the two variables used to

¹⁷El Niño refers to periodical variations in the Pacific water surface temperatures (MetOffice (2023)). Departures from average levels can cause worldwide consequences, affecting both weather and socio-economic outcomes (see Adams et al. (1999)).

capture weather events in this study: the SPEI and the cumulated precipitation anomalies.

Figure 1.4: Agro-ecological zones as per survey



Standardized Evapotranspiration Index (SPEI)

Climatologists consider droughts as multiscalar phenomena (see, for example, S. M. Vicente-Serrano, Beguería, and López-Moreno (2010)). Droughts can be broadly classified as Meteorological, Hydrological, Agricultural, or Socio-Economical. For this study, the most relevant one is the concept of agricultural drought, which, according to Hervás-Gámez and Delgado-Ramos (2019), can be thought of as “a period with declining soil moisture and consequent crop failure without any reference to surface water resources” (see also Rojas, Vrieling, and Rembold (2011)). Although droughts can all be defined as anomalies of precipitations with respect to long-run conditions leading to water deficits over a given period, they can differ by their onset, spatiotemporal extent, severity, and end (see Lloyd-Hughes (2014) and Mishra and Singh (2010)). More generally, meteorological water deficits are more challenging to assess than other natural phenomena that have been linked with economic outcomes (floods, for example; see Gröger and Zylberberg (2016)), and the use of rainfall to proxy for them is not always accurate, due to their high variability (Beguería, S. Vicente-Serrano, and Borja (2007)). A more important indicator, especially when studying agricultural phenomena, is soil moisture, a measure indicating the quantity of water retained by the ground (roughly speaking, drought is not necessarily caused by rainfall deficits as long as soil moisture remains relatively high).

Due to their multi-faceted nature, the quest for a single index summarizing all the defining aspects

of droughts has been challenging, and many indices have been developed for tracking different types of drought. The two most common are the Palmer Drought Severity Index (PDSI) and the Standardized Precipitation Index (SPI). The PDSI effectively tracks the elements determining water balance (prior precipitations, moisture, runoff and evaporation). The SPI allows to track water deficit accumulations over different time scales. Unfortunately, they both have shortcomings: the former is not flexible in terms of the time scale of the analysis (it is only available over 9-12 months), and the latter fails to take temperatures into account, which many studies found to be determinant for drought conditions.

An index effective in tracking the manifold aspects of droughts has been recently proposed by S. M. Vicente-Serrano, Beguería, and López-Moreno (2010): the Standardized Evapotranspiration Index (SPEI). In sum, the SPEI index combines precipitation and evapotranspiration (the amount of water released by the soil, which in turn depends on its composition, vegetation, and temperatures), providing a timely measure of the soil moisture content. The SPEI is an indicator centred around 0 and ranging approximately between -2 and $+2$, with positive values indicating wetter than normal period, while negative values indicate dryness and, eventually, drought. It can be calculated over different periods, allowing us to consider the cumulation of water deficits. This aspect, and the inclusion of temperature in constructing the index, makes the SPEI suitable for tracking different types of droughts. Moreover, the SPEI is a standardized variable which can be compared across time and space.

The conceptual novelty of the indicator is complemented by the recent release of a novel dataset by Peng et al. (2020). This dataset contains historical monthly values of the SPEI at different time scales for the African continent dating back to 1980 at a very fine resolution (approximately 5×55 km). In this work, I use 6-month and 12-month SPEI calculated in 4 different months (June, September, December, and March). Peng et al. (2020) suggest the following classification of climatic conditions based on a seven-values scale of the SPEI indicator, as portrayed in Table 1.7.

In particular, the preferred indicator adopted for the main empirical analysis is the 6-month SPEI in December. The reason for this choice depends on the characteristics of agricultural seasons in Ethiopia (as discussed at the beginning of this Section): most of the crop production is made in the Meher season, and the 6-month SPEI in December covers all the relevant phases (planting, growing, and harvesting, see Picture 1.3).

Table 1.7: SPEI values: categories

SPEI (s)	Meaning
$s \leq -2$	Extremely Dry
$-2 < s \leq -1.5$	Severely Dry
$-1.5 < s \leq -1$	Moderately Dry
$-1 < s - 1 < 1$	Near Normal
$1 \geq s < 1.5$	Moderately Wet
$1.5 \geq s < 2$	Very Wet
$s \geq 2$	Extremely Wet

The actual realizations of these indicators during the sample period are summarized in Table 1.8 (a cross-section summary of the values assumed by the different continuous indicators) and in Table 1.9, which reports the 6-month SPEI values in December based on the categorical distinction

made by Peng et al. (2020) and described in Table 1.7¹⁸. As a robustness exercise, I also create a further classification based on the SPEI values using three categories (more than dry, normal, and more than wet) by grouping extreme and severe values into a single indicator (more than dry and more than wet). The values based on this new classification are reported in Table 1.9. They show (consistently across the two specifications) how the first wave coincided with a relatively wet period and the second with a relatively normal one. In contrast, the third wave of the survey coincided with a very dry period (2015-2016 El Niño drought, as described by Negeri (2017b)).

Figure 1.5 reports analogous information by showing the shock distribution (6-month SPEI values measured in different months) across waves. It can be seen how the values of SPEI measured in June, September, and December are relatively similar (these indicators are the ones that matter for the agricultural season in the survey year), while SPEI measured in March is relatively different. Finally, another graphical representation of the 6-month SPEI relative to December is that of Figure 1.6: by noticing the legend at the right of each graph, it can be seen how the third wave has coincided with a particularly dry period¹⁹.

Table 1.8: Descriptives various continuous indicators

Variable	Mean	Std. Dev.	Min.	Max.	N
6-months SPEI in March	-0.06	0.99	-1.92	1.95	9493
6-months SPEI in June	-0.19	0.75	-2.31	2.07	9493
6-months SPEI in September	-0.01	1.08	-2.64	2.78	9493
6-months SPEI in December	0.07	1.14	-2.82	2.26	9493
12-months SPEI in March	-0.07	1.13	-2.64	2.22	9493
12-months SPEI in June	0.08	0.86	-2.31	2.18	9493
12-months SPEI in September	-0.28	0.93	-2.61	2.17	9493
12-months SPEI in December	-0.12	1.15	-2.9	2.9	9493


¹⁸In this table I only report 6-months SPEI because it is the preferred shock indicator in the empirical part of the study.

¹⁹The white areas in the graphs coincide with missing data. The lack of SPEI data is one of the reasons for dropping a few observations, as reported in Section 1.2.1

Table 1.9: Six months Spei distribution in December

	Wave			Total
	1	2	3	
Seven categories				
Extremely dry	0	0	630	630
Severely dry	12	12	684	708
Dry	108	32	345	485
Normal	2,295	1,874	1,465	5,634
Wet	228	896	116	1,240
Severely wet	378	312	8	698
Extremely wet	87	11	0	98
Total	3,108	3,137	3,248	9,493
Three categories				
More than dry	12	12	1,314	1,338
Normal	2,631	2,802	1,926	7,359
More than wet	465	323	8	796
Total	3,108	3,137	3,248	9,493

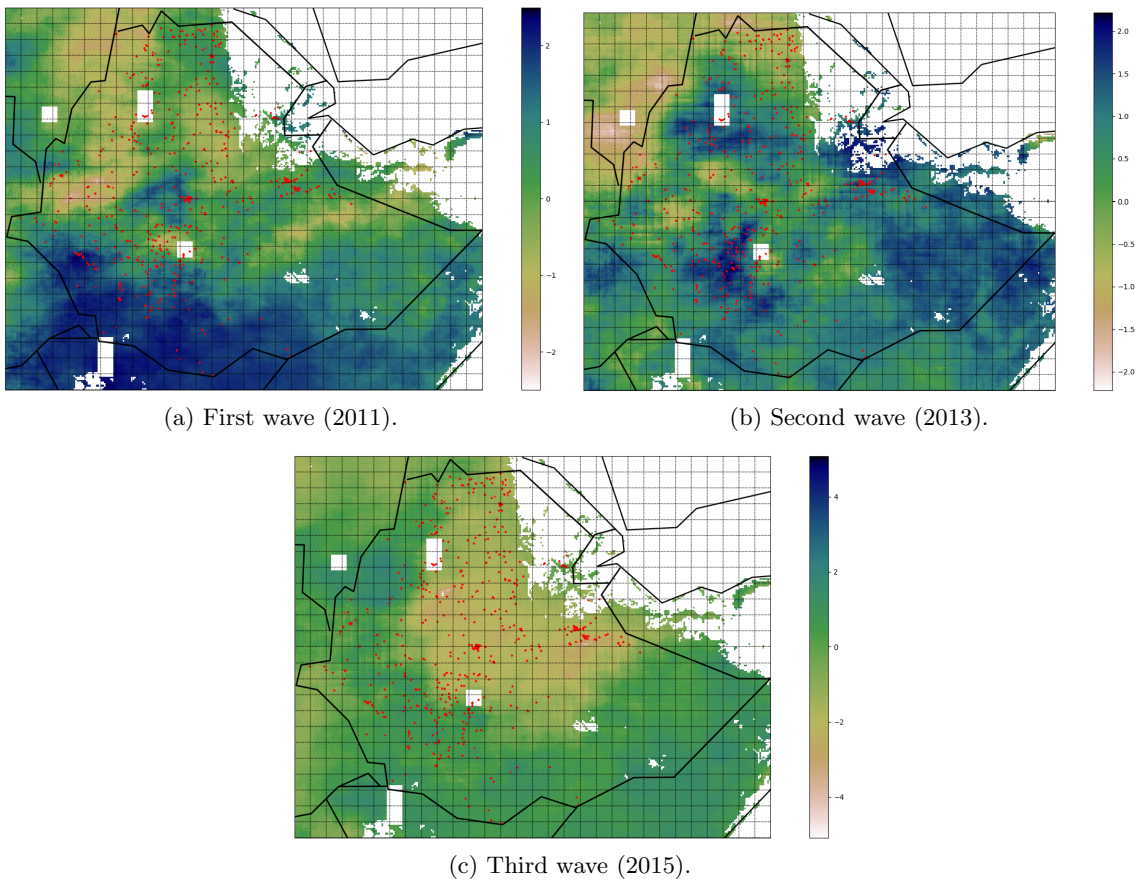
Figure 1.5: Distributions of Shocks across months of 6-months SPEI



Precipitations

Precipitation data are obtained from the Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS) dataset (Funk et al. (2015)). It provides rainfall time series with a 0.05° resolution, combining satellite imagery and in-situ station data. While my preferred indi-

Figure 1.6: 6-months SPEI values in December.



cator for weather shock is the SPEI index, I include precipitations in the analysis for robustness. Moreover, previous related studies on Ethiopian rural households (Dercon (2004) Demeke, Keil, and Zeller (2011), Porter (2012), Gao and Mills (2018)) relied exclusively on precipitations (although never, to the best of my knowledge, to satellite data with such a fine resolution - with the exception of Gao and Mills (2018)).

I built three precipitation indicators over different time spans for each location. The first one, which aims at capturing total rainfall in the Meher season, is the cumulated monthly rainfall between June and September. The second, to capture rainfall in the Belg season (the secondary rainy season in some parts of Ethiopia), is the cumulated monthly rainfall between March and April. The third one is the cumulated monthly rainfall between June and September, which aims at capturing the total amount of rainfall in all the months relevant to agriculture in the year of the surveys. For this last indicator, I plotted the long-run average value on the map of Ethiopia in Figure 1.7. This map, which is consistent with what is observed in the map of Figure 1.4, shows that households are scattered over a territory with a very diverse average climate (and rainfall, in particular).

The actual measure of the weather shocks is obtained by normalising the indicators described above by removing the long-run mean of the period (over the years preceding those of the surveys, i.e. 1981-2010) and dividing by the standard deviation. Such indicators are called meteorological anomalies and measure the distance of an actual year's precipitations in a given location from the long-run average in terms of numbers of standard deviations.

Figure 1.8 shows the different realisations of these indicators, and analogously to the map in Figure 1.7, the average level of cumulated precipitations in the three periods considered (bottom-right graph). Finally, Table 1.10 shows the various indicators' raw correlation (over the pooled sample). It can be noticed that the more traditional indicator of anomaly over the Meher season (June-August) is the least correlated with the other ones.

Figure 1.7: March-September long-run precipitation level (mm)

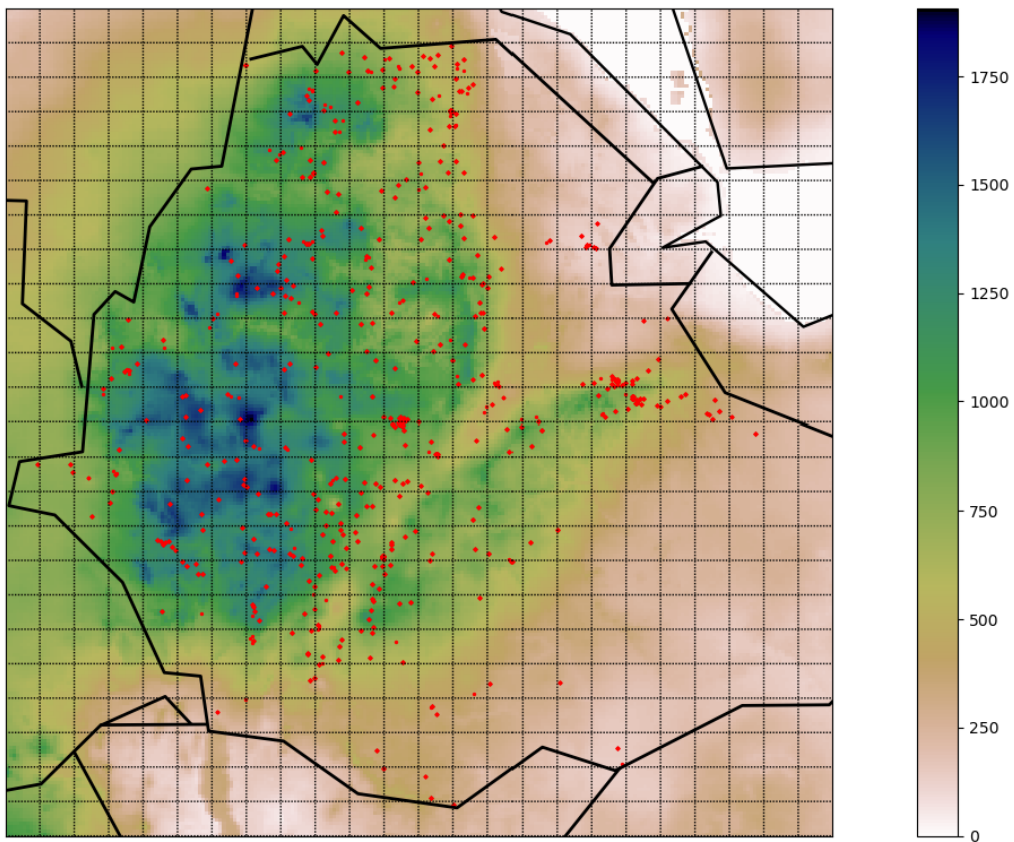


Figure 1.8: Precipitation anomalies and long-run levels

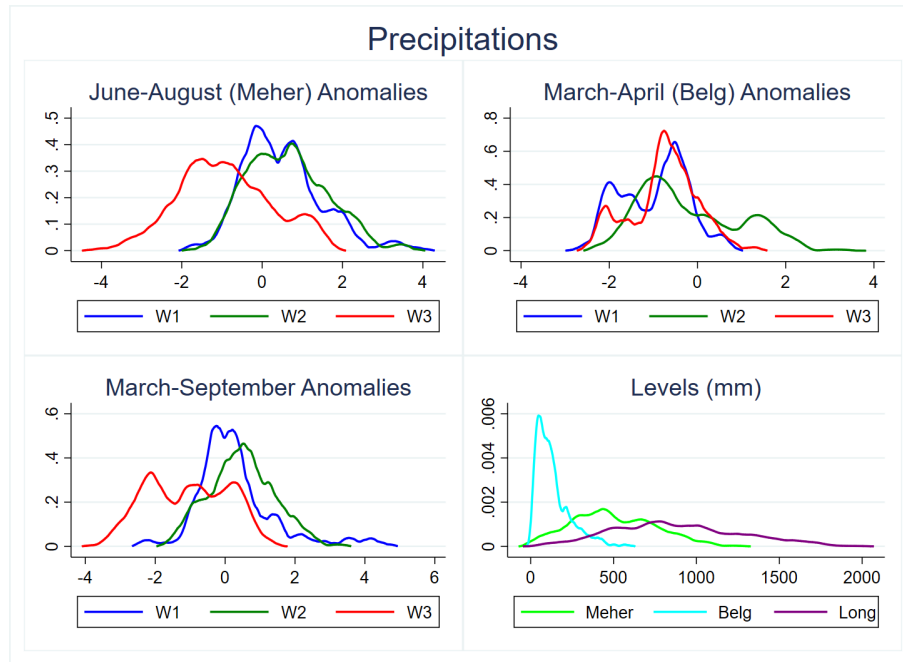


Table 1.10: Shocks correlations (across waves)

Variables	June-Aug anomaly	March-Aug anomaly	6-m spei september	6-m spei december	12-m spei september	12-m spei december
June-Aug anomaly	1.000					
March-Aug anomaly	0.715	1.000				
6-m spei september	0.777	0.963	1.000			
6-m spei december	0.740	0.917	0.947	1.000		
12-m spei september	0.698	0.967	0.968	0.925	1.000	
12-m spei december	0.697	0.953	0.971	0.959	0.968	1.000

1.2.3 The matching procedure

In this Section, I describe the procedure adopted to match the survey data, where households are geolocated through latitude-longitude coordinates, with the satellite weather dataset, provided as multidimensional data arrays. The procedure is implemented in Python, using tools contained in the standard libraries Numpy (Harris et al. (2020)) and Xarray (Hoyer and Hamman (2017))²⁰.

For every year, each latitude-longitude couple in the weather dataset can be thought of as a couple of integers determining their position (or order) in the data. For example, if latitude starts at 3°North and the grid has a precision of 0.05°, then the SPEI value at 3.15° is at node $\frac{3.15-3}{0.05} = 3$, i.e. at the third position in the grid. The analogous reasoning holds for longitudes.

The idea of the matching procedure is to associate each household location to the average of the four closest (sixteen as robustness) weather indicator measurements, as portrayed in Figure 1.9. The box below describes the various steps of the matching procedure through an example.

Example of the matching procedure (in a given wave t)

Steps in the algorithm.

1. Take Household i 's location, e.g. $(3.12^\circ N, 15.14^\circ E)$
2. Normalise it in terms of the satellite data grid $(\frac{15.14-15}{0.05}, \frac{3.12-3}{0.05}) = (2.8, 2.4)$
3. Take the closest nodes: $(2, 2)$, $(2, 3)$, $(3, 3)$, $(2, 3)$
4. Take the average of the weather indicator on those nodes
5. The household is associated with a weather value $SPEI_i = \sum_{j=1}^4 \frac{SPEI_j}{4}$

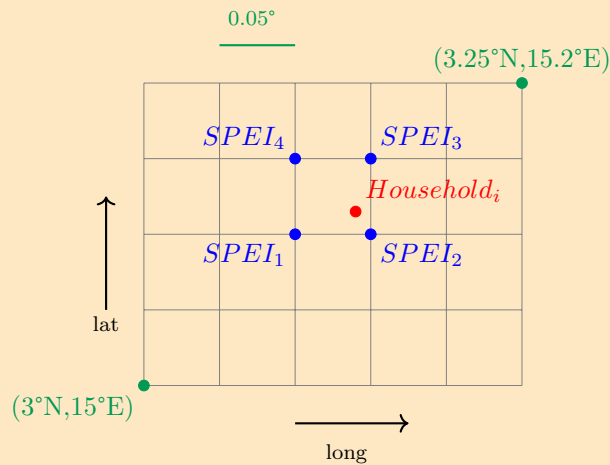


Figure 1.9: Matching procedure

²⁰The procedure is applied to both SPEI and precipitations (CHIRPS) data.

1.3 Theory and Empirical Specification

In this Section, I expose the theoretical rationale on which the analysis of the effects of shocks on consumption is based. I begin with a simple explanation of the classical intertemporal saving problem, which constitutes the natural point from where to start when analysing the ability of economic agents to insure from exogenous shocks. In the second part of this Section, I translate the theoretical model into its empirical counterpart, both for analysing the effects of shocks on consumption and to test the effectiveness of coping strategies.

1.3.1 Theoretical framework

There is a long theoretical tradition that seeks to explain what happens to consumption when the amount of available resources is subject to a change (for a review, look at Jappelli and Pistaferri (2010) and Meghir (2004)). Theory highlighted that consumption response depends on the nature of the shock: whether it is temporary or permanent on the one hand or whether it is anticipated or unexpected on the other. There are two main theoretical frameworks under which the ability of households to insure against shocks has been tested. The first is the self-insurance model or one of its variations; the second is the complete markets model. Both models give rise to testable predictions about consumption responses to income shocks.

In this section, mostly following Jappelli and Pistaferri (2017), I briefly outline the simplest version of the self-insurance model while I briefly expose and test the complete market model in the Appendix²¹. While the two frameworks are theoretically quite distinct, they are not always precisely distinguishable in their empirical specification, as pointed out by Bardhan and Udry (1999) and, previously, by Alderman and Paxson (1994).

The self-insurance model rests on an intertemporal utility maximisation problem: given its preferences, represented by the expected utility (1.1), the agent (or the household) seeks to choose streams of consumption $\{c_t\}_t$ given the period-by-period budget constraint (1.2) and the initial and terminal conditions (1.3):

$$\mathbb{E}_0 \sum_{t=0}^T \beta^t u(c_t), \quad (1.1)$$

and

$$a_{t+1} = (1 + r_{t+1})(a_t - c_t + y_t) \quad \text{for } t = 0, 1, \dots, T \quad (1.2)$$

$$\text{given } a_0 \quad \text{and} \quad a_{T+1} \geq 0. \quad (1.3)$$

The problem can be re-written in terms of dynamic programming, and the associated Bellman equation is:

$$V_t(a_t) = \max_{c_t, a_{t+1}} \{u(c_t) + \beta \mathbb{E}_t[V_{t+1}(a_{t+1})]\}; \quad (1.4)$$

²¹There are a few complications that can be added to the framework exposed here: the presence of liquidity, the precautionary savings motive, or a combination of the two. Although highly relevant, they are not necessary to expose the main features of the theory to be tested.

by inserting the budget constraint (1.2) in the right-hand side of (1.4), and by differentiating with respect to c_t and applying the Envelope theorem (see Stokey, Lucas, and Prescott (1989), Ljungqvist and Sargent (2012), or any other book on economic dynamics for the details, the procedure is relatively standard) it is possible to obtain the classical Euler equation, which states the condition that an optimal consumption stream must satisfy:

$$u'(c_t) = \beta \mathbb{E}_t[(1 + r_{t+1})u'(c_{t+1})]. \quad (1.5)$$

If one is willing to make a series of assumptions, the model boils down to the so-called certainty equivalence model, and the classical result by Hall (1978) is obtained, i.e. that consumption is a martingale: ²²

$$E_t[c_{t+1}] = c_t \quad \forall t \quad (1.6)$$

By defining the forecast error as $\varepsilon_{t+1} = c_{t+1} - E_t[c_{t+1}]$ and considering the intertemporal budget constraint at time t :

$$\sum_j^{T-t} \frac{E_t[c_{t+j}]}{(1+r)^j} = a_t \sum_j^{T-t} \frac{E_t[y_{t+j}]}{(1+r)^j}, \quad (1.7)$$

and letting $T \rightarrow +\infty$ we get that:

$$c_t = \frac{r}{1+r} \left[a_t + \sum_{j=0}^{+\infty} \frac{E_t[y_{t+j}]}{(1+r)^j} \right] \quad (1.8)$$

This equation tells the agent to set its consumption level to the annuity value of its total wealth, given by financial wealth plus the discounted value of expected income flow. Another expression for the sum of financial wealth and the agent's human capital is *permanent income*, from which the name Permanent Income Hypothesis often attributed to the certainty equivalent model. The same equation multiplied by $(1+r)$ becomes:

$$(1+r)c_t = r \left[a_t + \sum_{j=0}^{+\infty} \frac{E_t[y_{t+j}]}{(1+r)^j} \right] \quad (1.9)$$

while using the dynamic budget constraint (1.2) and forward by one period we obtain:

$$c_{t+1} = r(a_t + y_t - c_t) \sum_{j=0}^{+\infty} \frac{E_{t+1}[y_{t+j+1}]}{(1+r)^j} \quad (1.10)$$

Taking the difference between (1.10) and (1.9) we obtain another key prediction of the model:

$$\varepsilon_{t+1} = \Delta c_{t+1} = \frac{r}{1+r} \sum_{j=0}^{+\infty} \frac{E_{t+1}(y_{t+j+1}) - E_t(y_{t+j+1})}{(1+r)^j}, \quad (1.11)$$

²²The assumptions are: *i*) the constancy of the interest rate ($r_t = r \forall t$), *ii*) the equality of the discount rate and of the interest rate ($\beta(1+r) = 1$), *iii*) the linearity of marginal utility $u(c) = ac - \frac{b}{2}c_{t+1}^2$

which tells us that as long as there are no revisions in the agent's expectations also consumption will not change. Such restriction is at the basis of the modern tests of the Permanent Income Hypothesis. Overall, different models point out that economic agents use their savings or other forms of insurance to smooth their consumption across time periods. As a consequence, provided that there are no credit frictions or liquidity constraints, consumption shouldn't react much to anticipated variation in income.

As suggested by Jappelli and Pistaferri (2017), a relatively general empirical specification to test the sensitivity of consumption to changes in income, that reflects what the theory suggests, can be written as follows:

$$\Delta c_{it} = z'_{it}\lambda + \alpha E_{t-1}\Delta y_{it} + \sum_{j=1}^J \phi_j \pi_{it}^j + \xi_{it}, \quad (1.12)$$

where z_{it} is a set of controls for preference shifters such as age, family size or composition, $E_{t-1}\Delta y_{it}$ represents expected income changes, while π^j are different income shocks. Their coefficients represent the propensity of consume out of the different shocks. In the literature, there are two main approaches that such an equation allows to pursue: in the first one, the hypothesis that anticipated income changes do not affect consumption is tested ($\alpha = 0$); the second one tests the effect of unanticipated shocks on consumption, and it is based on the assumption that such shocks are exogenous (Jappelli and Pistaferri (2017) refer to one of the frameworks on which this approach is tested as the quasi-experimental approach). In this work, I follow the second approach and test the sensitivity of consumption of Ethiopian Rural households to various shocks: self-reported and objective ones (weather shocks). While I control for any shock that is recorded in the survey, the main focus is on the effect of objective shocks, those measured through satellite data on precipitations and the Standardised Precipitation-Evapotranspiration Index.

1.3.2 Empirical framework

In this work, I adopt two main empirical specifications, and this Section is devoted to their description. In the first part, in Section 1.3.2, I describe the model used to measure the effect of shocks on household consumption; in the second one, Section 1.3.2, I present a model that is an attempt of measuring the effectiveness of coping strategies adopted before or after the weather shock occurred.

Shocks and outcomes

In this section, I describe the empirical specification adopted to measure the response of household consumption to shocks. With respect to equation (1.12), I assume that $\alpha = 0$ and test the marginal propensity of consumption out of income shocks exploiting the within-household variation. I focus on the effect of shocks on consumption growth, and therefore the equation to be estimated is the following:

$$\log(c_{it}) = \mathbf{z}'_{it}\lambda + \phi \pi_{it} + \sum_{j=1}^J \sigma_j s_{it}^j + \mu_i + \delta D_t + \varepsilon_{it} \quad (1.13)$$

where \mathbf{z} is a vector of demographic and agricultural controls (see Table 1.1 and Table 1.2 for a description). In particular, following Dercon (2004), Gao and Mills (2018), and Porter (2012), the

vector \mathbf{z} includes the variable used as proxy for the Permanent Income, i.e. the number of Tropical Livestock Units the year previous to that of the survey. The variables $\{s_j\}_j$ are the indicators for self-reported shocks (see Table 1.3), while μ_i and D_t are, respectively the household and wave fixed-effect²³. Finally, ϕ represents the leading coefficient of interest, the effect of the objective weather shock on consumption. The weather shock is measured differently across specifications: it can be continuous (please refer to Table 1.8 or Figure 1.5 for a summary) or categorical (refer instead to Table 1.9 for a description of the discrete indicator in December). While I test different indicators, the preferred one (to be adopted also in analysing the effectiveness of coping strategies) is the 6-month SPEI index calculated in December.

I estimate the model using the *within estimator*, clustering the standard errors at the household level. In the section (1.B) of the Appendix, I also report the results for another specification, which focuses on the effects of shocks on consumption levels rather than logarithms, while in section 1.B.3 I report the estimation results when adopting the first differences of log-consumption (consumption growth), that are consistent with the results of the specification in levels.

Coping strategies effectiveness

To test the effectiveness of coping strategies, I add to the baseline model 1.13 the dummy variable r_{it} , which represents the adoption of a given coping strategy, and interact it with the shock index π_{it} . The various coping strategies are discussed in Section 1.2.1 and reported in Table 1.4²⁴. For the sake of clarity, I report a simplified version of the model (I omit the coefficients of self-reported shocks from Equation 1.14, but they are still present in the estimation phase) that allows us to focus on the coefficient of interest.

$$\log(c_{it}) = \mathbf{z}'_{it}\lambda + \phi\pi_{it} + \gamma r_{it} + \eta(\pi_{it} \times r_{it}) + \mu_i + \delta D_t + \varepsilon_{it} \quad (1.14)$$

The estimated coefficients $\hat{\gamma}$ and $\hat{\eta}$ measure, respectively, the effect on consumption of adopting a coping strategy and its effect in dampening (if actually effective) the weather shock. A completely effective coping strategy would make nihil the effect of the shock (i.e. the sum of $\hat{\phi}$ and $\hat{\eta}$ should be zero), a partially effective one would reduce the total effect of the shock ($\hat{\phi} + \hat{\eta} < \hat{\phi}$).

The estimating part of the exercise is analogous to that exposed in the previous subsection: the model includes time and household fixed effects, it is estimated through the *within estimator*, clustering the standard errors at the household level²⁵.

²³I do not explicitly separate shock into transitory or permanent.

²⁴The coping strategies tested are the presence of formal safety nets, the presence of informal safety nets, the participation to the PSNP, income diversification, the adoption of agricultural technologies, and the use of formal borrowing.

²⁵Using the standard Hausman test, the null hypothesis that the Random Effect estimator is the appropriate one is strongly rejected, for all the various specifications adopted.

1.4 Results

In this Section, I test the theory of the Permanent Income Hypothesis, and, in particular, the ability of Ethiopian rural households to insure their consumption against shocks. While, according to the data, households can insure against minor idiosyncratic shocks, the principal results tell a different story: on average, households cannot insure their consumption against weather shocks. This result is reasonable, especially in Ethiopia, where rural households heavily depend on agriculture. An exception is given when we group the values of the SPEI index into seven categories: with such a specification, households are able to insure consumption also against mild adverse weather shocks. A second interesting point is the presence of asymmetric effects, again when using the categorical indicator: mild positive shocks have a positive effect on household consumption, as found also by Porter (2012).

With the SPEI index, is it possible to accurately track the climatic conditions of the agricultural season, in particular by measuring the presence of climatic conditions dryer or wetter than normal ones. Combining the accuracy of the SPEI index with the dependency on agriculture and, therefore, on weather conditions, I quantitatively show the effects of weather shocks on the food consumption of rural households. The remainder of this Section is organised as follows: in Section 1.4.1 I first present the main results, showing how they are robust across the various specifications; in Section 1.4.2, I show the results when the sample is split in different ways (by income terciles, by long-run weather conditions, and by administrative region), and when the spatial correlation of the observations is taken into account; in Section 1.4.3 I report the evaluation of the effectiveness of coping strategies.

1.4.1 Shocks and outcome

This first Section of the results reports the estimates of the model linking directly shock and outcome. The next one will take into account the role of coping strategies.

Main results

In this Section, I report the main results of the estimation of the model expressed by Equation 1.13 by using both continuous indicators for weather shocks (see Table 1.11) and categorical ones (see Table 1.13). The models are estimated using the within estimator, clustering standard errors at the household level and including time and household fixed effects. The standard Hausman test rejects the null hypothesis that the random effect model is the correct specification.

The model controls for various demographic and agricultural indicators in order to get closer to the true effect of weather shocks on consumption (please refer to Tables 1.1 and 1.2 for the relative descriptives and Section 1.2.1 for a discussion). In particular, and consistently with the literature, livestock holdings are positively correlated with consumption, while the dependency ratio is negatively correlated. Other coefficients have the expected sign, negative if the family head is female and positive if the head is literate, but are not significantly different from zero. In section 1.4.2, several robustness exercises are conducted and discussed, while in Appendix I report the results of estimations when using the level of food consumption instead of the logarithm and the results of estimation when allowing for spatial correlation of the observations. The results remain substantially unchanged across all specifications.

Turning to the main object of the study, the effects on consumption of weather shocks, all continuous indicators show a positive coefficient across the columns of Table 1.11: an increase in cumulative precipitations, as captured by the first two columns, and an increase in the SPEI index

measured over different periods and in different moments of the year, are both positively correlated with the logarithm of households consumption. Although qualitatively similar, as expected after looking at correlations between indicators in Table 1.10, some specifications capture higher effects than others. In particular, the stronger effect is associated with the use of the 6-month SPEI measured in December, while the lower one is associated with the use of cumulative precipitations during the rainy (Meher) season. These two results are sensible, for the 6-month SPEI measured in December captures the overall conditions in the agricultural season preceding the household interviews.

With respect to the stronger effect, an increase (decrease) of the indicator by one unit increases (decreases) consumption by about 6%. Polynomial specifications of degrees two and three (not shown but available) do not substantially change the results. If we accept such a linear relationship, households hit by the strongest shock in the sample (for example, during the drought that hit the country during the third wave of the Survey) lost up to 20% of consumption. Overall, even using the indicator delivering the smallest magnitude, such as the anomaly of cumulative precipitations across June, July, and August (first column of Table 1.11), can have a sizeable effect on consumption, especially if we think that the household in the sample are often close to the poverty line. This, in turn, highlights that rural households, on average, cannot insure consumption against major weather shocks. From an economic viewpoint, this is not surprising if we think that weather shocks can be considered an example of aggregate shocks (in contrast with idiosyncratic ones), which are harder to insure.

To further highlight this issue, in Table 1.12 I report the coefficient of the variables indicating self-reported shocks (these variables are included as controls in the main model presented in Table 1.11, I report them separately to make tables easier to read). All the coefficients of such variables are not significantly different from zero. This result might be an indication that households are actually able to insure against several types of shocks, but it might be due to lack of variation in the data (see Table 1.3). As a further robustness check, in Table 1.18, I report the results of the estimation of the baseline model with continuous indicator when grouping the self-reported shocks indicator into two categories, aggregate and idiosyncratic shocks, as explained in Section 1.2.1 and in Footnote 8, and reported in Table 1.3. It can be seen that even when grouping the shocks, the results are stable, and the coefficients of these two new variables are not statistically different from zero, hinting again at a certain ability of households to insure against minor shocks.²⁶

In order to explore possible asymmetries of the effect of positive and negative shocks on consumption, such as the ones reported by Porter (2012) and Gao and Mills (2018), in Table 1.13 I report the results obtained by using categorical indicators for weather shocks: most indicators only capture the effect of negative shocks (e.g. experiencing a more than dry agricultural season), consistently with Porter (2012). The only exception arises when looking at the 6-month SPEI measured in December, which captures a positive effect when households experience a wetter-than-normal agricultural season. In particular, a dry season is associated with an overall decrease in food consumption by about 18%, while a wet season increases food consumption by about 14 percentage points. In terms of magnitudes, the results obtained by using the categorical version of the indicator are consistent with the results obtained using the continuous one: a more than dry season, by construction, means a SPEI indicator below 2 (reaching values of almost -3 in the data, see Table 1.8), and therefore it is delivering analogous results to those obtained using the continuous indicator.

²⁶One caveat that must be added to this discussion, in addition to that on the possible lack of variation in the data, is that being shocks self-reported, there might be some measurement error (self-reporting bias); to corroborate this view, I noticed that a self-reported indicator of drought is only weakly correlated with objective shocks indicator such SPEI or precipitations.

Finally, in Table 1.14, I report the results of grouping the SPEI values into seven categories. Interestingly, and again in line with the results obtained by Porter (2012), Table 1.14 suggests that households are able to insure against mild weather shocks. When using a finer grouping of the values of the SPEI index (as suggested by Peng et al. (2020)), it can be seen that experiencing a (mildly) dry season has no significant effect on consumption.

Table 1.11: Effect on consumption of weather shocks: different continuous indicators

dep. var: (log) food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
Dep. Ratio (non working / working)	-0.047***	-0.048***	-0.047***	-0.046***	-0.048***	-0.047***
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Household head is female	-0.064	-0.062	-0.063	-0.062	-0.064	-0.062
	(0.044)	(0.044)	(0.044)	(0.044)	(0.044)	(0.044)
Household's head age	0.015**	0.015**	0.015**	0.014**	0.015**	0.014**
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Household's head age squared	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Household's head can write and read	0.031	0.033	0.032	0.034	0.033	0.034
	(0.025)	(0.025)	(0.025)	(0.025)	(0.025)	(0.025)
Household Size	0.072***	0.072***	0.071***	0.072***	0.072***	0.072***
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Household mean age	-0.009***	-0.009***	-0.009***	-0.009***	-0.009***	-0.009***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
TLU 1 y. ago	0.013***	0.013***	0.013***	0.013***	0.013***	0.013***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
continuous indicator	0.024***	0.028***	0.037***	0.063***	0.031***	0.049***
	(0.008)	(0.007)	(0.009)	(0.008)	(0.009)	(0.008)
constant	5.851***	5.862***	5.881***	5.907***	5.867***	5.895***
	(0.427)	(0.427)	(0.427)	(0.425)	(0.430)	(0.427)
R-squared overall	0.174	0.166	0.168	0.166	0.168	0.157
F	10.265	10.273	10.488	12.095	10.383	11.169
N	9493	9493	9493	9493	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level
s.e. in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Author's own calculations

Further specifications

In Section 1.A of the Appendix, I report the results when using the SPEI indicator calculated also in March. In this case, the indicator is not predictive of food consumption. This result appears sensible since both the 6-month and the 12-month indicators refer to a period too far from that investigated during the survey. This result is also consistent with the fact that controlling for the first meteorological season occurring between two waves (for example, if we consider the third wave, there is the meteorological season in t , occurring between February and September 2015, and there is the meteorological season in $t - 1$, occurring between February and September 2014; see Figure 1.2 for a clarification) does not affect the results.

In Section 1.B of Appendix I report the estimation results of the specification that uses the level of food consumption rather than the logarithm. While qualitatively the results are the same, quantitatively they are interesting because they help in picturing the effect of the shocks in terms of a monetary metric (2016 US Dollars). In Table 1.23, it can be seen that depending on the indicator, an increase (decrease) of the indicator by one unit implies an increase (decrease) of consumption between 17 and 75 2016 US Dollars.

In Tables 1.24 and 1.25, I report the results in levels when using the categorical indicators (grouping the SPEI values, respectively, in three and seven categories, as suggested by Peng et al. (2020)). Again, the results are qualitatively in line with the specification that uses the logarithm of food consumption.

Table 1.12: Coefficients of self reported shocks in the main specification (see Table 1.11)

dep. var: (log) food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
death bread winner-HEAD	-0.052 (0.053)	-0.052 (0.053)	-0.052 (0.053)	-0.049 (0.053)	-0.052 (0.053)	-0.052 (0.053)
illness Household member	0.024 (0.022)	0.023 (0.022)	0.022 (0.022)	0.025 (0.022)	0.024 (0.022)	0.025 (0.022)
loss non farm job Household member	-0.100 (0.065)	-0.099 (0.065)	-0.098 (0.065)	-0.092 (0.066)	-0.103 (0.065)	-0.094 (0.065)
landslides/avalanches	-0.044 (0.093)	-0.049 (0.093)	-0.047 (0.093)	-0.036 (0.093)	-0.050 (0.093)	-0.042 (0.094)
heavy rain preventing work	0.038 (0.053)	0.039 (0.053)	0.040 (0.053)	0.022 (0.053)	0.041 (0.053)	0.023 (0.053)
other crop damage	-0.021 (0.032)	-0.021 (0.032)	-0.018 (0.032)	-0.018 (0.032)	-0.020 (0.032)	-0.022 (0.032)
price fall of food item	0.010 (0.039)	0.010 (0.039)	0.011 (0.039)	0.022 (0.038)	0.012 (0.039)	0.015 (0.038)
price raise of food item	0.027 (0.021)	0.029 (0.021)	0.029 (0.021)	0.033 (0.020)	0.030 (0.021)	0.030 (0.021)
increase in price of inputs	-0.025 (0.025)	-0.028 (0.025)	-0.028 (0.025)	-0.035 (0.025)	-0.028 (0.025)	-0.033 (0.025)
great loss/death livestock	-0.015 (0.031)	-0.005 (0.031)	-0.007 (0.031)	-0.015 (0.031)	-0.006 (0.031)	-0.012 (0.031)
fire	-0.198 (0.123)	-0.199 (0.124)	-0.198 (0.124)	-0.189 (0.124)	-0.199 (0.124)	-0.194 (0.123)
theft/robbery/violence	0.090 (0.080)	0.090 (0.080)	0.088 (0.079)	0.088 (0.080)	0.090 (0.080)	0.089 (0.080)
loss of house/farm/land	-0.111 (0.107)	-0.109 (0.108)	-0.112 (0.107)	-0.110 (0.107)	-0.113 (0.107)	-0.111 (0.107)
displacement (govt. proj)	-0.109 (0.380)	-0.115 (0.376)	-0.120 (0.376)	-0.117 (0.375)	-0.111 (0.377)	-0.120 (0.374)
local unrest/violence	-0.225*** (0.087)	-0.217** (0.086)	-0.213** (0.086)	-0.228*** (0.084)	-0.221** (0.086)	-0.223*** (0.085)

Notes. This table reports the coefficients of the self-reported shocks variable as included in the baseline model of Table 1.11

s.e. in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Author's own calculations

1.4.2 Robustness and Heterogeneity

In this Section, I report the results when controlling for different household characteristics. I split the sample according to several features: by farm income terciles, by agro-zones, by long-run precipitations, and by administrative regions. Such subdivisions are made to check the robustness of the results and, importantly, to investigate whether different socio-economic characteristics can influence how shocks affect household welfare. Finally, I take into account the spatial structure of the data by considering different geographical aggregations of the weather variables, by including lagged weather indicators to account for longer horizon climatic conditions, and by assuming different levels of spatial autocorrelation in the errors. The results remain robust to such specifications. All the robustness checks are reported in the Appendix (see Section 1.A).

Farm income

Farm income is the results of sales of agricultural products (including livestock) minus expenditures (all intermediates and inputs: seeds, fertilizers, pesticides, transportation costs, labour, capital, and land renting) plus any agricultural wage income. In this, I follow Magalhães and Santaaulàlia-Llopis (2018), which also discuss the possibility of income under-reporting in household surveys. While farm income seems to proxy quite well consumption in autarkic rural households, I prefer to rely only on the position (wave by wave) of the households in the income distribution rather

Table 1.13: Effect on consumption of weather shocks: different categorical indicators

dep. var: (log) food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
more than dry	0.009 (0.025)	-0.162*** (0.024)	-0.149*** (0.025)	-0.182*** (0.025)	-0.161*** (0.024)	-0.176*** (0.026)
more than wet	0.029 (0.022)	0.003 (0.026)	0.039 (0.027)	0.140*** (0.027)	0.097 (0.062)	0.043 (0.033)
constant	5.809*** (0.429)	5.944*** (0.427)	5.912*** (0.423)	5.921*** (0.414)	5.901*** (0.426)	5.948*** (0.421)
R-squared overall	0.173	0.183	0.181	0.162	0.174	0.179
F	9.405	11.204	10.912	11.850	11.054	11.319
N	9493	9493	9493	9493	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level
s.e. in parentheses, * p<0.10, ** p<0.05, *** p<0.01

Source: Author's own calculations

Table 1.14: Effect on consumption of weather shocks: categorical Spei in different months

dep. var: (log) food	6m-Spei mar	6m-Spei jun	6m-Spei sep	6m-Spei dec	12m-Spei mar	12m-Spei jun	12m-Spei sep	12m-Spei dec
ext. dry	-0.047 (0.063)	-0.081 (0.104)	-0.123*** (0.044)	-0.197*** (0.033)	0.211* (0.121)	0.018 (0.055)	-0.142*** (0.035)	-0.183*** (0.030)
sev. dry	0.135*** (0.025)	0.099** (0.041)	-0.162*** (0.028)	-0.179*** (0.032)	-0.116* (0.060)	-0.032 (0.046)	-0.178*** (0.029)	-0.189*** (0.036)
dry	0.007 (0.020)	0.049** (0.021)	-0.020 (0.028)	-0.038 (0.030)	0.045 (0.031)	0.029 (0.029)	-0.030 (0.024)	-0.038 (0.026)
normal	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)
wet	-0.021 (0.026)	-0.045 (0.030)	-0.048 (0.030)	0.007 (0.022)	-0.078*** (0.020)	-0.123*** (0.020)	-0.094*** (0.029)	0.036* (0.022)
sev. wet	0.194** (0.088)	0.140*** (0.054)	-0.023 (0.032)	0.136*** (0.028)	-0.004 (0.033)	-0.111** (0.043)	0.103 (0.076)	-0.015 (0.037)
ext. wet		0.153 (0.161)	0.140*** (0.048)	0.189** (0.075)	-0.235 (0.173)	-0.672*** (0.181)	0.053 (0.099)	0.213*** (0.061)
constant	5.848*** (0.424)	5.770*** (0.436)	5.899*** (0.418)	5.936*** (0.416)	5.794*** (0.425)	5.778*** (0.424)	5.908*** (0.425)	5.945*** (0.427)
R-squared overall	0.187	0.187	0.187	0.160	0.177	0.176	0.185	0.169
F	9.326	9.020	9.908	10.553	9.213	9.793	10.197	10.410
N	9493	9493	9493	9493	9493	9493	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level
s.e. in parentheses, * p<0.10, ** p<0.05, *** p<0.01

Source: Author's own calculations

than using the precise income value for other exercises²⁷. I group the sample households in income terciles and test the effect of weather shocks across them (captured, as usual, by the 6-month SPEI in December). These results are reported in the first column of Table 1.19.

The results are reassuring, at least about the data quality. While the asymmetric effects of the shock are confirmed (negative for conditions dryer than normal and positive for wetter conditions), the average consumption is increasing in the income terciles (although significantly only for the third tercile). More importantly, while the negative shock hits the second and third deciles harder than the first one, the third is hit less hard than the second. This result may hint to a higher ability of richer households to insure food consumption²⁸.

²⁷The point is that while I do not entirely trust the point value of the reported income, I trust that income data are at least able to capture the relative position of households in the income distribution

²⁸These results must be taken with a grain of salt: income data are contemporaneous to consumption one and might be affected by weather shocks in the same way (and in fact, this is quite likely). I keep using such a robustness check because I rely on relatively broad quantiles of the income distribution, and, being the weather

Agro-zones

I discussed the presence of very different agro-ecological zones in Ethiopia in Section 1.2.2. Based on the distinct agro-zones into which household locations are classified, I group them into two newly defined categories: arid or non-arid agro-ecological zones. The second column of Table 1.19 reports the robustness checks of splitting the sample into these two categories. The relevant point in this case is that negative weather shocks seem to hit harder zones that are already arid. The rationale for such a robustness check was to investigate whether households in more arid zones were somehow more prepared against adverse shocks (droughts) due to the development of long-run adaptation strategies. This is not the case, at least according to the data, and in fact, arid zones suffer more than non-arid in case of adverse shocks.

Long-run precipitations

A further check, similar in spirit to that exposed in the previous section, is based on splitting the sample on the base of the long-run precipitation average. Again, I split the sample into three terciles, looking at the long-run cumulative precipitation average distribution across March to September. The idea is to check whether being in a more or less arid zone favours households against adverse shocks. The most notable result in this case is that being in an arid zone, i.e. in the dryer tercile, accentuates the effects of both negative and positive shocks.

Administrative regions

This last robustness check, in terms of sample-splitting, was done primarily for completeness, and I did not expect anything relevant to emerge from it. I divided the sample according to the official administrative regions of the country to which household locations belong. While the general results that, on average, the households do not insure against strong adverse weather shocks holds, it is harder to interpret the coefficients of the interactions. This result might be due to the fact that even administrative regions are extremely diverse (see Oromya, for example), and no clear pattern seems to emerge.²⁹

Spatial autocorrelation

As a final robustness check, I take into account the spatial structure of the data through three different exercises, reported in Section 1.A.2 of the Appendix. In the first one, reported in Table 1.11, I re-estimate the baseline model by assuming different levels of spatial autocorrelation in the errors. I do it using the Stata command *acreg*, developed by Colella et al. (2023). The authors developed an estimator for the variance-covariance matrix that allows for arbitrary structure in the correlation between the error terms in regression analysis. In particular, such an estimator can be used in the presence of spatial and time correlation. This situation is particularly likely to happen in contexts similar to the one depicted in this study, i.e., in the presence of weather shocks and geographically clustered observations (e.g. Enumerations Areas or villages). The command uses latitude-longitude coordinates and the time structure of the sample to correct for eventual biases. One option allows setting the threshold below which observations are to be considered correlated. I test several values of such threshold (50, 100, 200 Kms), but in Table 1.20, I only report the results when using a threshold of 50 Kms. The Table shows that the results hold across indicators, although standard errors are inflated, and some continuous indicators show a slight

shock an aggregate one, it should not affect the relative order in the distribution.

²⁹The results of the estimation are not reported but available upon request.

drop in their significance level. The preferred indicator (6-month SPEI measured in December) remains significant at the 1% level.

In the second exercise, I aggregate the continuous weather indicators across all households belonging to the same area (I report the results of aggregating the weather indicator over a grid of $5^\circ \times 5^\circ$ in Table 1.21 to exclude that the household-specific weather indicators might capture some household characteristics that vary over time but that is not observable and that co-vary with weather. The results of such an exercise are consistent, both qualitatively and in terms of the magnitude of the relevant coefficients, with the estimates of the baseline model reported in Table 1.11.

In the third exercise, I control for the weather conditions one year before the survey. In this way, I control for weather shocks contemporaneous to the interview and the conditions in the previous season (see 1.2 for a graphical representation). Moreover, I also consider the interaction between the indicator in t and in $t - 1$. The idea is to control for the potential increased vulnerability due to repeated bad seasons. I use the 12-month SPEI indicator registered in December of the year before the survey to capture the overall dryness condition in the year not covered in the survey. Table 1.22 shows that the lagged weather indicator is not significantly different from zero but that the interaction coefficient is positive and significant. This result suggests a reinforcing effect of bad (or good) seasons of food consumption.

1.4.3 Coping strategies

In this Section, I report the results on the effectiveness of coping strategies. As mentioned previously (see Section 1.2.1), some coping strategies are not mutually exclusive, and therefore I prefer to test them separately. I use both the continuous and categorical indicators to measure the weather shock, and I report the results in the columns, respectively, of Table 1.15 and Table 1.16. I describe, for each of the two specifications, the results for each coping strategy, which are reported in the columns of the two tables. Since the results are analogous when food is measured in levels rather than in logarithms, I limit the discussion to the latter specification and only report the results of the former in Tables 1.26 and 1.27 in Section 1.B.2 of the Appendix.

Overall, only two strategies can be defined as effective: the possibility to smooth income through occupations different from the main agricultural one and the possibility to borrow from formal institutions (e.g., banks). Unfortunately, the data do not allow to go beyond statistical associations and plausible stories, and precise causal effects remain unproven (while the shocks considered in the previous Section can be considered exogenous, too many selection mechanisms, such as those described in the following discussion, are at play).

Formal safety nets

The test of this first coping strategy is reported in the first column of Tables 1.15 and 1.16. As a reminder, the variable representing the coping strategy indicates the presence of any transfer from Governmental and non-governmental organisations. Both Tables consistently show that using such coping strategies is associated with a worsening in the effect of the shock with respect to the average. This might be due to the fact that interventions of Governmental and non-governmental organisations are triggered and targeted where conditions are already associated with vulnerability. This result is partly consistent with the findings by Gao and Mills (2018), who also find a (stronger) effectiveness of formal safety nets in smoothing consumption against weather shocks.

Informal safety nets

This second coping strategy is associated with households receiving any gift or transfer (cash, food, or in-kind) from relatives or friends. The second column of Tables 1.15 and 1.16 tells a slightly different story than the one of formal safety nets. Neither the coping strategy nor the interaction coefficient are significantly different from zero. This result is reasonable if we think that these sorts of transfers can be used against idiosyncratic shocks, where the shock hits one individual (household), but the relative or the friend is not. In the presence of weather shocks, which are aggregate shocks, this type of insurance becomes harder to achieve, and the empirical results are consistent with it.

Income smoothing

Having more than one source of income allows households to insure against weather shocks. Column three of Tables 1.15 and 1.16 reports the effects of having family members with different occupations from the main agricultural one. While such a condition alone does not increase the average consumption level, it dampens the overall effect of the adverse shocks. This result is partly consistent with the findings by Gao and Mills (2018), who also find a (milder) effectiveness of off-farm labour in smoothing consumption against weather shocks.

Household in PSNP

The story and the quantitative results for the participation in the PSNP program are similar to those of having some kind of formal safety net (in fact, the two conditions partly overlap, as discussed previously). As it can be seen in the fourth column of Tables 1.15 and 1.16, the difference is that participating in the PSNP actually increases the average consumption of the households, but the program is not effective in dampening the shock. This result might be due again to a selection mechanism: PSNP targets particularly vulnerable households, particularly prone and susceptible to shocks.

Agricultural technologies adoption

This coping strategy refers to the use of improved seeds, the use of fertilisers, or the cultivation of several types of cropping. I refer to these practices as ex-ante coping strategies because, in principle, they should help rural households become more resilient to shocks. The fifth column of Tables 1.15 and 1.16 shows that while adopting certain technologies is associated with higher consumption, it has no effect in softening the effect of weather shock. The positive effect may be due once more to a selection mechanism: due to education, economies of scale, or other reasons, the adoption of better technologies might be mainly relegated to wealthier families.

Formal borrowing

The last columns of Tables 1.15 and 1.16 show that having the possibility to resort to formal borrowing is not only associated with higher consumption but almost nullifies the effect of the adverse weather shock. While the former effect might be due to the selection mechanism at play (as discussed in the discussion of the other coping strategies), the latter might actually indicate what the theory predicts. In the presence of well-functioning credit markets, households borrow in order to smooth consumption.

Table 1.15: Different coping strategies, continuous shock indicators

dep. var: (log) food	formal transf.	informal transf.	off-farm labour	PSNP	ex-ante agric.	credit
6-months SPEI in dec	0.057*** (0.009)	0.060*** (0.009)	0.072*** (0.009)	0.059*** (0.009)	0.047*** (0.014)	0.073*** (0.009)
coping strat.	-0.003 (0.021)	0.015 (0.021)	0.032 (0.029)	0.087*** (0.030)	0.051** (0.026)	0.048** (0.023)
coping strat. × 6-months SPEI in dec	0.025* (0.014)	0.014 (0.017)	-0.073*** (0.019)	0.032* (0.019)	0.019 (0.015)	-0.071*** (0.018)
constant	5.903*** (0.425)	5.896*** (0.425)	5.886*** (0.426)	5.899*** (0.424)	5.866*** (0.428)	5.846*** (0.431)
R-squared overall	0.164	0.165	0.150	0.161	0.160	0.137
F	11.703	11.460	11.753	11.670	11.611	12.033
N	9493	9493	9493	9488	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level
s.e. in parentheses, * p<0.10, ** p<0.05, *** p<0.01

Source: Author's own calculations

Table 1.16: Different coping strategies, categorical shock indicators

dep. var: (log) food	formal transf.	informal transf.	off-farm labour	PSNP	ex-ante agric.	credit
MT dry	-0.143*** (0.030)	-0.184*** (0.026)	-0.206*** (0.026)	-0.157*** (0.027)	-0.182*** (0.047)	-0.218*** (0.027)
MT wet	0.129*** (0.030)	0.138*** (0.029)	0.163*** (0.028)	0.163*** (0.027)	0.142** (0.055)	0.141*** (0.028)
coping strat.	0.011 (0.023)	0.013 (0.023)	0.018 (0.032)	0.146*** (0.032)	0.054** (0.027)	0.024 (0.026)
MT dry × coping strat.	-0.127*** (0.043)	0.017 (0.058)	0.198*** (0.061)	-0.178*** (0.054)	0.001 (0.049)	0.207*** (0.054)
MT wet × coping strat.	0.067 (0.062)	0.008 (0.073)	-0.181** (0.080)	-0.256** (0.105)	-0.001 (0.062)	-0.016 (0.109)
constant	5.924*** (0.414)	5.911*** (0.414)	5.895*** (0.419)	5.890*** (0.410)	5.887*** (0.414)	5.865*** (0.421)
R-squared overall	0.155	0.162	0.135	0.153	0.162	0.127
F	11.695	10.997	11.234	11.874	10.985	11.658
N	9493	9493	9493	9488	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level
s.e. in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Author's own calculations

1.5 Conclusions

Households in developing countries face a substantial amount of risk, which affects their welfare and vulnerability to poverty. While households exhibit a certain ability to insure against shocks, previous studies showed that they are far from reaching the benchmark of perfect insurance, both from the intertemporal viewpoint and in terms of within-community risk sharing. I focused on rural households in Ethiopia, a country historically prone to weather shocks and which, despite growing, remains very poor. This work touches on two macro-areas of economics research. The first one is about risk and coping strategies in developing countries and, from a theoretical viewpoint, on consumption smoothing across time and states of nature; the second one is a relatively new strand of research called climate economics, which systematically tests the effects of weather and climate on economic outcomes.

I contribute to the literature by merging these two branches of research and applying climate economics tools (the use of satellite data) on household panel data to test consumption insurance theories and the ability of rural households in Ethiopia to deal with risk. In a certain sense, one of the contributions consists in providing new answers to relatively old questions: on the one hand, the test of consumption theories gave rise to an extensive literature that is still relevant nowadays. On the other hand, the possibility of relying on new and more precise data improves the quality of the answers.

This work provides two types of evidence of the ability of rural households to smooth consumption against shocks. On the one hand, households seem able to insure against most idiosyncratic and mild adverse weather shocks. On the other hand, vulnerability to stronger weather shocks (especially droughts) remains elevated. Further evidence shows that specific coping strategies (the possibility to smooth income through occupations different from the main agricultural one and the possibility to borrow from formal institutions) can alleviate the effects even of strong weather shocks.

A secondary result obtained through this work is the proof of the suitability of satellite data, and

in particular of data regarding the SPEI index, to test consumption insurance theories and, more generally, to track household welfare in rural Ethiopia. Incidentally, this work is also one of the firsts using the SPEI index in economics (to the best of my knowledge there are only other three: Azzarri and Signorelli (2020), Albert, Bustos, and Ponticelli (2021), and Piolatto et al. (2022)). This ability to track weather conditions timely can be very relevant in a shock-prone country such as Ethiopia, especially under the additional risk of climate change.

1.A Robustness

In this Section, I report the estimates of the model when using the SPEI indicator calculated also in March (Table 1.17), the coefficients of the baseline model estimated once self-reported shocks are grouped into two categories -idiosyncratic vs. aggregate- (Table 1.18), and the results of the various robustness checks explained in Section 1.4.2.

1.A.1 Robustness and heterogeneity

Table 1.17: Effect on consumption of weather shocks: continuous Spei in different months

dep. var: (log) food	6m-Spei mar	6m-Spei jun	6m-Spei sep	6m-Spei dec	12m-Spei mar	12m-Spei jun	12m-Spei sep	12m-Spei dec
indic	-0.007 (0.011)	0.010 (0.010)	0.037*** (0.009)	0.063*** (0.008)	0.002 (0.008)	-0.040*** (0.008)	0.031*** (0.009)	0.049*** (0.008)
R-squared overall	0.180	0.173	0.168	0.166	0.177	0.182	0.168	0.157
F	9.641	9.697	10.488	12.095	9.680	10.311	10.383	11.169
N	9493	9493	9493	9493	9493	9493	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level
s.e. in parentheses, * p<0.10, ** p<0.05, *** p<0.01

Source: Author's own calculations

Table 1.18: Coefficients of the model estimated by grouping self-reported shocks into two categories.

dep. var: (log) food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
idiosyncratic shock	-0.016 (0.019)	-0.018 (0.019)	-0.018 (0.019)	-0.016 (0.019)	-0.018 (0.019)	-0.016 (0.019)
aggregate shock	-0.025 (0.016)	-0.021 (0.016)	-0.020 (0.016)	-0.019 (0.016)	-0.019 (0.017)	-0.021 (0.016)
TLU 1 y. ago	0.013*** (0.003)	0.013*** (0.003)	0.013*** (0.003)	0.013*** (0.003)	0.013*** (0.003)	0.013*** (0.003)
continuous indicator	0.024*** (0.008)	0.028*** (0.007)	0.038*** (0.009)	0.061*** (0.008)	0.029*** (0.009)	0.048*** (0.008)
constant	5.822*** (0.422)	5.831*** (0.422)	5.850*** (0.422)	5.872*** (0.420)	5.832*** (0.425)	5.863*** (0.422)
R-squared overall	0.146	0.139	0.141	0.139	0.141	0.131
F	18.159	18.207	18.624	21.416	18.349	19.865

Notes. This table reports the coefficients of the self-reported shocks grouped into two variables as explained in Section 1.2.1 and Footnote 8 s.e. in parentheses, * p<0.10, ** p<0.05, *** p<0.01

Source: Author's own calculations

1.A.2 Spatial features

Table 1.19: Heterogeneity analysis: farm income terciles, agro-zones, long-run precipitations, and administrative regions

dep. var: (log) food	income terciles	agro-zones	long run prec.
more than dry	-0.101** (0.044)	-0.076** (0.034)	-0.188*** (0.035)
more than wet	0.138*** (0.051)	0.142*** (0.029)	0.103** (0.047)
2nd farm income tercile	0.038 (0.024)		
3rd farm income tercile	0.097*** (0.025)		
more than dry × 2nd farm income tercile	-0.139** (0.056)		
more than dry × 3rd farm income tercile	-0.117** (0.050)		
more than wet × 2nd farm income tercile	-0.024 (0.065)		
more than wet × 3rd farm income tercile	0.014 (0.068)		
arid agro-zone		-0.007 (0.249)	
more than dry × arid agro-zone		-0.199*** (0.039)	
more than wet × arid agro-zone		-0.044 (0.065)	
more than dry × low rain			-0.100** (0.043)
more than dry × high rain			0.179*** (0.052)
more than wet × low rain			0.106 (0.066)
more than wet × high rain			0.011 (0.062)
R-squared overall	0.006	0.003	0.002
F	15.684	22.133	20.333
N	9493	9493	

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level
s.e. in parentheses, * p<0.10, ** p<0.05, *** p<0.01

Source: Author's own calculations

1.B Other specifications

In this section I report the effects of weather shocks on consumption measured in levels rather than in logarithms. Since the model specification and other considerations are analogous to those exposed in section 1.3.2 of the main text, I will only briefly report them in the following subsections for the sake of completeness.

1.B.1 Shocks and outcomes

I estimate the model:

$$c_{it} = \mathbf{z}'_{it}\lambda + \phi\pi_{it} + \sum_{j=1}^J \sigma_j s^j_{it} + \mu_i + \delta D_t + \varepsilon_{it}. \quad (1.15)$$

The only difference with the model exposed in section 1.3.2 is that consumption is measured in 2016 US Dollars rather than taking the logarithm. All the remaining details about the specification and estimation are essentially unchanged.

Table 1.20: Effect on consumption of weather shocks: continuous indicators and spatial correlation

dep. var: (log) food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
continuous indicator	0.024*	0.028**	0.037**	0.063***	0.031*	0.049***
	(0.014)	(0.013)	(0.016)	(0.015)	(0.016)	(0.014)
R-squared	0.037	0.038	0.038	0.044	0.037	0.041
N	9493	9493	9493	9493	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are corrected to account for spatial autocorrelation using households' locations
s.e. in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Author's own calculations

Table 1.21: Effect on consumption of weather shocks: continuous indicators at a higher aggregation

dep. var: (log) food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
continuous indicator	0.025***	0.033***	0.042***	0.062***	0.038***	0.053***
	(0.008)	(0.007)	(0.009)	(0.008)	(0.009)	(0.009)
R-squared overall	0.136	0.124	0.127	0.130	0.127	0.119
F	18.026	18.395	18.709	21.109	18.637	19.983
N	9493.000	9493.000	9493.000	9493.000	9493.000	9493.000

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; the weather indicator is aggregated (averaged) at a higher geographical level ($0.5 \times 0.5^\circ$)
s.e. in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Author's own calculations

1.B.2 Coping strategies

1.B.3 First differences

Table 1.22: Effect on consumption of weather shocks: continuous indicators and lags

dep. var: (log) food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
continuous indicator	0.022*** (0.008)	0.025*** (0.007)	0.031*** (0.009)	0.051*** (0.008)	0.025*** (0.009)	0.040*** (0.009)
lg	0.006 (0.008)	0.013 (0.008)	0.007 (0.008)	0.008 (0.008)	0.007 (0.009)	0.015* (0.008)
c.indic × c.lg	0.013** (0.006)	0.022*** (0.006)	0.025*** (0.007)	0.033*** (0.007)	0.011 (0.009)	0.026*** (0.008)
R-squared overall	0.136	0.121	0.126	0.121	0.127	0.109
F	16.495	17.164	17.423	20.202	16.448	18.500
N	9493.000	9493.000	9493.000	9493.000	9493.000	9493.000

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; these specifications control for weather conditions in the year previous to the survey, and the interaction with current year indicator

s.e. in parentheses, * p<0.10, ** p<0.05, *** p<0.01

Source: Author's own calculations

Table 1.23: Effect on consumption of weather shocks: levels

dep. var: food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
Dep. Ratio (non working / working)	-46.362*** (9.896)	-46.664*** (9.837)	-46.405*** (9.825)	-44.938*** (9.739)	-46.661*** (9.837)	-46.271*** (9.803)
Household head is female	-75.009* (41.839)	-71.996* (41.684)	-72.609* (41.782)	-72.198* (41.734)	-73.411* (41.616)	-71.756* (41.638)
Household's head age	11.959** (6.072)	11.895** (6.050)	11.476* (6.029)	11.296* (6.055)	11.296* (6.045)	11.344* (6.055)
Household's head age squared	-0.085 (0.054)	-0.084 (0.054)	-0.081 (0.054)	-0.080 (0.054)	-0.079 (0.054)	-0.080 (0.054)
Household's head can write and read	-3.124 (30.519)	-0.885 (30.420)	-1.923 (30.422)	0.064 (30.082)	-1.105 (30.459)	0.064 (30.354)
Household Size	45.360*** (10.840)	45.336*** (10.764)	44.236*** (10.778)	45.681*** (10.664)	45.505*** (10.755)	45.744*** (10.707)
Household mean age	-5.408*** (1.669)	-5.347*** (1.660)	-5.370*** (1.661)	-5.127*** (1.653)	-5.307*** (1.662)	-5.237*** (1.660)
TLU 1 y. ago	15.710*** (4.655)	15.261*** (4.625)	15.215*** (4.635)	15.438*** (4.517)	15.383*** (4.637)	15.412*** (4.578)
continuous indicator	17.544** (7.155)	33.312*** (7.136)	47.097*** (8.837)	75.328*** (8.591)	41.793*** (8.584)	53.095*** (8.304)
constant	450.243 (350.240)	479.023 (348.086)	506.556 (348.088)	533.756 (347.799)	494.017 (350.023)	509.332 (348.717)
R-squared overall	0.112	0.109	0.108	0.108	0.110	0.109
F	7.321	7.557	7.801	8.855	7.663	8.058
N	9493	9493	9493	9493	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level

s.e. in parentheses, * p<0.10, ** p<0.05, *** p<0.01

Source: Author's own calculations

Table 1.24: Effect on consumption of weather shocks: levels

dep. var: food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
more than dry	45.630*	-193.112***	-184.140***	-224.101***	-222.537***	-211.485***
	(23.351)	(25.387)	(26.863)	(26.573)	(26.864)	(26.941)
more than wet	-5.867	-44.877**	-11.269	94.783***	22.482	20.183
	(19.444)	(22.759)	(23.880)	(23.759)	(56.800)	(27.128)
constant	390.539	591.527*	552.042	566.079*	555.803	587.774*
	(352.660)	(349.861)	(345.519)	(338.091)	(345.237)	(341.807)
R-squared overall	0.112	0.101	0.101	0.102	0.098	0.099
F	6.756	8.552	8.148	8.586	8.671	8.363
N	9493	9493	9493	9493	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level

s.e. in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Author's own calculations

Table 1.25: Effect on consumption of weather shocks: levels

dep. var: food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
ext. dry	98.874***	-224.412***	-238.012***	-291.405***	-267.171***	-263.196***
	(31.122)	(31.318)	(57.452)	(38.587)	(48.669)	(33.323)
sev. dry	-10.312	-135.583***	-167.330***	-170.520***	-192.872***	-152.074***
	(28.237)	(34.283)	(29.296)	(31.770)	(29.012)	(36.594)
dry	-14.334	-22.712	-7.834	-17.881	-8.545	-25.907
	(34.872)	(27.510)	(26.464)	(28.347)	(22.778)	(27.080)
wet	-0.723	-18.934	-29.419	-3.521	-96.935***	-4.053
	(23.398)	(26.668)	(22.804)	(22.122)	(22.641)	(22.981)
sev. wet	-38.227	-67.638**	-45.531	103.306***	14.534	-25.148
	(25.941)	(28.638)	(28.905)	(26.333)	(77.729)	(30.256)
ext. wet	40.412	-26.466	35.866	32.077	-5.202	122.051**
	(28.153)	(32.048)	(43.550)	(48.348)	(66.216)	(53.482)
constant	384.314	588.260*	550.507	552.193	540.489	556.461
	(348.139)	(347.568)	(343.658)	(339.331)	(345.012)	(346.745)
R-squared overall	0.109	0.100	0.099	0.101	0.095	0.100
F	6.463	7.634	7.364	7.657	8.339	7.781
N	9493	9493	9493	9493	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level

s.e. in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Author's own calculations

Table 1.26: Different coping strategies, continuous shock indicators

dep. var: food	formal transf.	informal transf.	off-farm labour	PSNP	ex-ante agric.	credit
6-months SPEI in dec	71.129***	75.681***	85.523***	70.656***	53.536***	87.817***
	(9.484)	(8.836)	(9.215)	(8.475)	(12.856)	(9.455)
coping strat.	-15.238	-5.754	0.725	88.068**	55.624**	25.112
	(24.047)	(24.446)	(37.966)	(34.785)	(25.409)	(23.101)
coping strat. \times 6-months SPEI in dec	17.944	-1.791	-83.675***	43.281	27.685*	-83.405***
	(15.681)	(21.245)	(19.349)	(28.121)	(14.753)	(19.317)
constant	534.478	537.146	507.670	525.882	486.405	473.003
	(347.629)	(347.927)	(346.911)	(346.214)	(351.818)	(351.743)
R-squared overall	0.109	0.108	0.109	0.119	0.106	0.114
F	8.535	8.394	8.744	8.468	8.413	8.893
N	9493	9493	9493	9488	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level

s.e. in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Source: Author's own calculations

Table 1.27: Different coping strategies, categorical shock indicators

dep. var: food	formal transf.	informal transf.	off-farm labour	PSNP	ex-ante agric.	credit
MT dry	-190.026*** (32.188)	-231.016*** (27.963)	-252.657*** (27.513)	-205.103*** (27.591)	-248.583*** (45.041)	-273.541*** (28.992)
MT wet	91.511*** (25.818)	102.692*** (26.359)	112.684*** (25.709)	107.392*** (24.733)	50.877 (47.626)	92.567*** (24.463)
1.indic	2.827 (27.697)	-7.484 (25.472)	-22.928 (42.479)	131.902*** (37.534)	47.439* (26.985)	-14.683 (26.400)
MT dry × 1.indic	-108.692** (51.267)	40.077 (73.881)	241.686*** (79.821)	-139.940* (74.820)	32.889 (51.535)	285.983*** (64.628)
MT wet × 1.indic	20.768 (58.587)	-48.915 (58.106)	-137.644** (62.088)	-142.072* (81.531)	55.926 (53.089)	27.587 (100.415)
constant	570.665* (337.191)	572.454* (340.462)	523.470 (339.893)	544.965 (335.314)	544.327 (340.905)	505.295 (345.109)
R-squared overall	0.110	0.103	0.108	0.117	0.100	0.109
F	8.321	7.989	8.454	8.159	8.111	8.598
N	9493	9493	9493	9488	9493	9493

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level
s.e. in parentheses, * p<0.10, ** p<0.05, *** p<0.01

Source: Author's own calculations

Table 1.28: Effect on consumption of weather shocks: first differences

dep. var: food	Jun-Aug anomaly	Mar-Aug anomaly	6m-Spei sep	6m-Spei dec	12m-Spei sep	12m-Spei dec
Dep. Ratio (non working / working)	-0.028*** (0.009)	-0.028*** (0.009)	-0.028*** (0.009)	-0.027*** (0.009)	-0.028*** (0.009)	-0.028*** (0.009)
Household head is female	0.033 (0.021)	0.033 (0.021)	0.033 (0.021)	0.035* (0.021)	0.032 (0.021)	0.034 (0.021)
Household's head age	-0.008** (0.004)	-0.008** (0.004)	-0.008** (0.004)	-0.008** (0.004)	-0.008** (0.004)	-0.008** (0.004)
Household's head age squared	0.000** (0.000)	0.000** (0.000)	0.000*** (0.000)	0.000*** (0.000)	0.000*** (0.000)	0.000*** (0.000)
Household's head can write and read	0.010 (0.019)	0.011 (0.019)	0.011 (0.019)	0.012 (0.019)	0.011 (0.019)	0.013 (0.019)
Household Size	-0.008* (0.005)	-0.008 (0.005)	-0.008* (0.005)	-0.007 (0.005)	-0.007 (0.005)	-0.007 (0.005)
Household mean age	-0.007*** (0.001)	-0.007*** (0.001)	-0.007*** (0.001)	-0.006*** (0.001)	-0.007*** (0.001)	-0.007*** (0.001)
id_shock	-0.026 (0.024)	-0.029 (0.024)	-0.030 (0.024)	-0.029 (0.024)	-0.028 (0.024)	-0.029 (0.024)
aggr_shock	-0.062*** (0.020)	-0.052** (0.020)	-0.050** (0.020)	-0.049** (0.020)	-0.051** (0.020)	-0.052** (0.020)
TLU 1 y. ago	0.007*** (0.002)	0.007*** (0.002)	0.007*** (0.002)	0.007*** (0.002)	0.007*** (0.002)	0.007*** (0.002)
Household uses formal financial services	0.142*** (0.027)	0.143*** (0.027)	0.143*** (0.027)	0.146*** (0.027)	0.143*** (0.027)	0.144*** (0.027)
continuous indicator	0.050*** (0.010)	0.047*** (0.009)	0.065*** (0.012)	0.067*** (0.011)	0.050*** (0.011)	0.055*** (0.010)
R-squared overall						
F	11.710	12.248	12.472	12.383	11.734	12.193
N	6130.000	6130.000	6130.000	6130.000	6130.000	6130.000

Notes. self-reported shocks, agricultural, and geographical variables are used but not reported; household and time FE are present; errors are clustered at the household level
s.e. in parentheses, * p<0.10, ** p<0.05, *** p<0.01

Source: Author's own calculations

References

- Adams, RM et al. (1999). “The economic consequences of ENSO events for agriculture”. In: *Climate Research* 13, pp. 165–172.
- Ahmed, Musa Hasen, Wondimagegn Mesfin Tesfaye, and Franziska Gassmann (Feb. 2023). “Early growing season weather variation, expectation formation and agricultural land allocation decisions in Ethiopia”. In: *Journal of Agricultural Economics* 74 (1), pp. 255–272.
- Albert, Christoph, Paula Bustos, and Jacopo Ponticelli (July 2021). *The Effects of Climate Change on Labor and Capital Reallocation*.
- Alderman, Harold and Christina H. Paxson (1994). “Do the Poor Insure? A Synthesis of the Literature on Risk and Consumption in Developing Countries”. In: Palgrave Macmillan UK, pp. 48–78.
- Attanasio, Orazio (1999). “Chapter 11 Consumption”. In: pp. 741–812.
- (Mar. 2015). “Frank Ramsey’s a Mathematical Theory of Saving”. In: *The Economic Journal* 125 (583), pp. 269–294.
- Attanasio, Orazio and Guglielmo Weber (Sept. 2010). “Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy”. In: *Journal of Economic Literature* 48 (3), pp. 693–751.
- Auci, Sabrina, Laura Castellucci, and Manuela Coromaldi (2018). “The impact of climate change on the distribution of rural income in Ethiopia”. In: *International Journal of Environmental Studies* 75.6, pp. 913–931.
- Auci, Sabrina and Manuela Coromaldi (2021). “Climate variability and agricultural production efficiency: evidence from Ethiopian farmers”. In: *International Journal of Environmental Studies* 78.1, pp. 57–76.
- Auffhammer, Maximilian et al. (2013). *Using Weather Data and Climate Model Output in Economic Analyses of Climate Change*.
- Azzarri, Carlo and Sara Signorelli (Jan. 2020). “Climate and poverty in Africa South of the Sahara”. In: *World Development* 125, p. 104691.
- Bardhan, Pranab and Christopher Udry (1999). *Development microeconomics*. OUP Oxford.
- Beguéría, Santiago, Sergio Vicente-Serrano, and Reig Fergus Borja Latorre and (2007). *About the SPEI Drought monitoring and analysis based on climatic indices*. URL: <https://spei.csic.es/home.html>.
- Bodkin, Ronald (1959). “Windfall Income and Consumption”. In: *The American Economic Review* 49.4, pp. 602–614.
- Campbell, John Y (Nov. 1987). “Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis”. In: *Econometrica* 55.6, pp. 1249–1273.
- Carleton, Tamma A. and Solomon M. Hsiang (Sept. 2016). “Social and economic impacts of climate”. In: *Science* 353 (6304).
- Carter, Michael R. and Travis J. Lybbert (2012). “Consumption versus asset smoothing: Testing the implications of poverty trap theory in Burkina Faso”. In: *Journal of Development Economics* 99 (2), pp. 255–264.
- Colella, Fabrizio et al. (2023). “acreg: Arbitrary correlation regression”. In: *The Stata Journal* 23.1, pp. 119–147.
- Colmer, Jonathan (May 2021). “Rainfall Variability, Child Labor, and Human Capital Accumulation in Rural Ethiopia”. In: *American Journal of Agricultural Economics* 103 (3), pp. 858–877.
- Deaton, Angus (Oct. 1992). *Understanding Consumption*. Oxford University Press Oxford.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (2014). “learnfromweather”. In: *Journal of Economic Literature* 52 (3), pp. 740–798.

- Demeke, Abera Birhanu, Alwin Keil, and Manfred Zeller (Sept. 2011). “Using panel data to estimate the effect of rainfall shocks on smallholders food security and vulnerability in rural Ethiopia”. In: *Climatic Change* 108 (1), pp. 185–206.
- Dercon, Stefan (2004). “Growth and shocks: Evidence from rural Ethiopia”. In: *Journal of Development Economics* 74 (2), pp. 309–329.
- Dercon, Stefan and Luc Christiaensen (2011). “Consumption risk, technology adoption and poverty traps: Evidence from Ethiopia”. In: *Journal of Development Economics* 96 (2), pp. 159–173.
- Dercon, Stefan, John Hoddinott, and Tassew Woldehanna (2005). “Shocks and consumption in 15 Ethiopian villages, 1999–2004”. In: *Journal of African Economies* 14 (4), pp. 559–585.
- Deschênes, Olivier and Michael Greenstone (Mar. 2007). “The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather”. In: *American Economic Review* 97.1, pp. 354–385.
- FAO (2023). *About the SPEI Drought monitoring and analysis based on climatic indices*. URL: <https://www.fao.org/giews/countrybrief/country.jsp?code=ETH&lang=ar>.
- Flavin, Marjorie A. (Oct. 1981). “The Adjustment of Consumption to Changing Expectations About Future Income”. In: *Journal of Political Economy* 89 (5), pp. 974–1009.
- Friedman, Milton (May 1957). *A Theory of the Consumption Function*. NBER Books frie57-1. National Bureau of Economic Research, Inc.
- Funk, Chris et al. (Dec. 2015). “The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes”. In: *Scientific Data* 2 (1), p. 150066.
- Gao, Jianfeng and Bradford F. Mills (2018). “Weather Shocks, Coping Strategies, and Consumption Dynamics in Rural Ethiopia”. In: *World Development* 101, pp. 268–283.
- Gröger, André and Yanos Zylberberg (Apr. 2016). “Internal Labor Migration as a Shock Coping Strategy: Evidence from a Typhoon”. In: *American Economic Journal: Applied Economics* 8.2, pp. 123–53.
- Hall, Robert E. (Dec. 1978). “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence”. In: *Journal of Political Economy* 86 (6), pp. 971–987.
- Hall, Robert E. and Frederic S. Mishkin (Mar. 1982). “The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households”. In: *Econometrica* 50.2, pp. 461–481.
- Harris, Charles R. et al. (Sept. 2020). “Array programming with NumPy”. In: *Nature* 585.7825, pp. 357–362.
- Hervás-Gámez, Carmen and Fernando Delgado-Ramos (2019). “Drought management planning policy: from Europe to Spain”. In: *Sustainability* 11.7, p. 1862.
- Hill, Ruth Vargas and Catherine Porter (Aug. 2017). “Vulnerability to Drought and Food Price Shocks: Evidence from Ethiopia”. In: *World Development* 96, pp. 65–77.
- Hoyer, S. and J. Hamman (2017). “xarray: N-D labeled arrays and datasets in Python”. In: *Journal of Open Research Software* 5.1.
- Hsiang, Solomon (Oct. 2016). “Climate econometrics”. In: *Annual Review of Resource Economics* 8 (1), pp. 43–75.
- Hsiang, Solomon M. and Kyle C. Meng (May 2015). “Tropical economics”. In: vol. 105. American Economic Association, pp. 257–261.
- Jappelli, Tullio and Luigi Pistaferri (2010). “The Consumption Response to Income Changes”. In: *Annual Review of Economics* 2 (1), pp. 479–506.
- (Oct. 2017). *The Economics of Consumption*. Vol. 1. Oxford University Press.
- Kazianga, Harounan and Christopher Udry (2006). “Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso”. In: *Journal of Development Economics* 79 (2), pp. 413–446.

- Kinnan, Cynthia (2022). “Distinguishing Barriers to Insurance in Thai Villages”. In: *Journal of Human Resources* 57.1, pp. 44–78.
- Kochar, Anjini (1999). “Smoothing Consumption by Smoothing Income: Hours-of-Work Responses to Idiosyncratic Agricultural Shocks in Rural India”. In: *The Review of Economics and Statistics* 81.1, pp. 50–61.
- Kocherlakota, N. R. (Oct. 1996). “Implications of Efficient Risk Sharing without Commitment”. In: *The Review of Economic Studies* 63 (4), pp. 595–609.
- Ligon, Ethan (1998). “Risk Sharing and Information in Village Economies”. In: *The Review of Economic Studies* 65.4, pp. 847–864.
- Ljungqvist, Lars and Thomas J. Sargent (Feb. 2012). *Recursive Macroeconomic Theory, Third Edition*. Vol. 1. MIT Press Books 0262018748. The MIT Press.
- Lloyd-Hughes, Benjamin (2014). “The impracticality of a universal drought definition”. In: *Theoretical and Applied Climatology* 117 (3-4), pp. 607–611.
- LSMS-ISA and CSAoE (Dec. 2012). “Living Standards Measurement Study- Integrated Surveys on Agriculture (LSMS-ISA): Ethiopia Rural Socioeconomic Survey (ERSS) - Basic Information Document”. In.
- (Mar. 2015). “Ethiopia Socioeconomic Survey (ESS) Wave Two (2013/2014) - Basic Information Document”. In.
- (Feb. 2017). “Ethiopia Socioeconomic Survey (ESS) Wave Three (2015/2016) - Basic Information Document”. In.
- Magalhães, Leandro De and Raül Santaaulàlia-Llopis (2018). “The consumption, income, and wealth of the poorest: An empirical analysis of economic inequality in rural and urban Sub-Saharan Africa for macroeconomists”. In: *Journal of Development Economics* 134, pp. 350–371.
- Manski, Charles F. (Sept. 2004). “Measuring Expectations”. In: *Econometrica* 72 (5), pp. 1329–1376.
- Meghir, Costas (2004). “A Retrospective on Friedman’s Theory of Permanent Income”. In: *The Economic Journal* 114.496, F293–F306.
- MetOffice (2023). *El Niño, La Niña and the Southern Oscillation*. URL: <https://www.metoffice.gov.uk/research/climate/seasonal-to-decadal/gpc-outlooks/el-nino-la-nina/enso-description>.
- Mishra, Ashok K. and Vijay P. Singh (Sept. 2010). “A review of drought concepts”. In: *Journal of Hydrology* 391 (1-2), pp. 202–216.
- Modigliani, Franco and Richard Brumberg (Feb. 1954). *The Collected Papers of Franco Modigliani, Volume 6*. Vol. 6. MIT Press Books 0262134543. The MIT Press.
- Morduch, Jonathan (Sept. 1995). “Income Smoothing and Consumption Smoothing”. In: *Journal of Economic Perspectives* 9.3, pp. 103–114.
- Negeri, Melkamu Belina (2017a). “The Effects of El Nino on Agricultural GDP of Ethiopia”. In: *American Journal of Water Science and Engineering* 3 (4), p. 45.
- (2017b). “The Effects of El Nino on Agricultural GDP of Ethiopia”. In: *American Journal of Water Science and Engineering* 3 (4), p. 45.
- Paxson, Christina H. (1992). “Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand”. In: *The American Economic Review* 82.1, pp. 15–33.
- (Feb. 1993). “Consumption and Income Seasonality in Thailand”. In: *Journal of Political Economy* 101 (1), pp. 39–72.
- Peng, Jian et al. (Mar. 2020). “A pan-African high-resolution drought index dataset”. In: *Earth System Science Data* 12 (1), pp. 753–769.
- Piolatto, Amedeo et al. (2022). *New insights on distributive politics from a new way to measure aridity*. eng.

- Porter, Catherine (2012). “Shocks, Consumption and Income Diversification in Rural Ethiopia”. In: *Journal of Development Studies* 48 (9), pp. 1209–1222.
- Ramsey, F. P. (Dec. 1928). “A Mathematical Theory of Saving”. In: *The Economic Journal* 38 (152), p. 543.
- Rojas, O., A. Vrieling, and F. Rembold (Feb. 2011). “Assessing drought probability for agricultural areas in Africa with coarse resolution remote sensing imagery”. In: *Remote Sensing of Environment* 115 (2), pp. 343–352.
- Rosenzweig, Mark R and Kenneth I Wolpin (1993). *Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India*, pp. 223–244.
- Sebastian, Kate (2009). *Agro-ecological Zones of Africa*. Version DRAFT VERSION. URL: <https://doi.org/10.7910/DVN/HJYYTI>.
- Skoufias, Emmanuel and Agnes R Quisumbing (Mar. 2005). “Consumption Insurance and Vulnerability to Poverty: A Synthesis of the Evidence from Bangladesh, Ethiopia, Mali, Mexico and Russia”. In: *The European Journal of Development Research* 17 (1), pp. 24–58.
- Stokey, Nancy L., Robert E. Lucas, and Edward C. Prescott (Oct. 1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.
- Townsend, Robert M. (1994). “Risk and Insurance in Village India”. In: *Econometrica* 62.3, pp. 539–591.
- (1995). “Consumption Insurance : An Evaluation of Risk-Bearing Systems in Low-Income Economies”. In: *Journal of Economic Perspectives* 9 (3), pp. 83–102.
- United Nations (2012). *Relief Web - A Climate Trend Analysis of Ethiopia*. URL: <https://reliefweb.int/report/ethiopia/climate-trend-analysis-ethiopia?gclid=Cj0KCQjw0bunBhD9ARIsAAZ1OE3wCB>.
- Vicente-Serrano, Sergio M., Santiago Beguería, and Juan I. López-Moreno (Apr. 2010). “A multiscalar drought index sensitive to global warming: The standardized precipitation evapotranspiration index”. In: *Journal of Climate* 23 (7), pp. 1696–1718.
- Wakjira, Mosisa Tujuba et al. (Nov. 2021). “Rainfall seasonality and timing: implications for cereal crop production in Ethiopia”. In: *Agricultural and Forest Meteorology* 310, p. 108633.
- Wolpin, Kenneth I (1982). *A New Test of the Permanent Income Hypothesis: The Impact of Weather on the Income and Consumption of Farm Households in India*, pp. 583–594.

Chapter 2

Trends of concentration in selected OECD countries

2.1 Introduction

This chapter is part of a larger project conducted in the Productivity Innovation and Entrepreneurship Division of the STI Directorate at the OECD¹. The work aims at studying the evolution of several proxies of competition. There is a growing literature documenting recent trends on market concentration, markups, entrenchment, and other measures that can proxy trends in competition. Understanding the evolution of such trends is important for policy reasons but also in the context of macroeconomics. Measuring trends can help understand what generated them and also what they might affect.

Studies on concentration differ in their definition of a market – both in terms of the product or industry scope and the geographical coverage – and in their specification of firm boundaries across markets and countries, such as through the role of business groups.

This study makes numerous contributions to the literature on competition in OECD countries. It covers in an integrated way multiple measures that proxy for different facets of competition (or a lack of it): concentration and entrenchment. Examining different characteristics of markets in the same setting provides richer inferences about the trends in market power, as well as possible explanations for these trends. The analysis is conducted for both manufacturing and services (and beyond, as it also includes mining and utilities sectors) across many OECD countries and over a long-time horizon. Constructing a database that allows such analysis is itself a significant contribution.

The project innovates on the existing literature in its measurement of concentration, aimed at reflecting markets more accurately. First, it measures concentration within more narrowly defined industries than most previous cross-country studies, mainly at the 3-digit level. Second, international trade is accounted for by constructing a taxonomy that defines whether markets compete domestically or internationally and computes concentration at the corresponding geographic level. Therefore, industries are classified into three geographical buckets: domestic, European, or global.

¹This project is a joint work with Christian Abele, Sara Calligaris, Chiara Criscuolo, Josh De Lyon, Andrea Greppi, and Miguel Chaves.

Using this taxonomy means that the concentration measures account for firms' international activities, even when firm-level trade data is absent. On top of this, in the robustness checks, imports and exports are incorporated in the concentration measure using industry-level data. Third, following the methodology by Bajgar, Berlingieri, et al. (2019), the connectedness of firms within business groups is accounted for to incorporate the role of mergers and acquisitions in driving concentration trends and to capture the complete activities of multinational firms in a market.

Alongside concentration, a static measure of market shares at any point in time, the project also measures entrenchment. Entrenchment is a dynamic measure of the persistence of firms as market leaders and provides richer insights into the extent of competition, even when concentration is high. The measurement of entrenchment innovates on previous literature by defining markets tradeable internationally following the taxonomy and by accounting for the connectedness of firms within a business group.

In terms of empirical results, on average, industry concentration has increased across all geographic buckets. Industries that compete at the domestic level had the greatest increase in average concentration, by around 6 percentage points (p.p., henceforth) between 2000 and 2019. Industries that compete internationally – either at the European or global level – increased their concentration by approximately 4 p.p.

The comparison of unweighted and weighted concentration cumulative changes shows that when weighting for the relative importance of the markets, concentration looks overall flat over the period 2000-2019 for industries competing both domestically and at the European level. In industries competing at the global level, the weighted trend is even decreasing. The combination of these results suggests that for the domestic and European buckets, the increase in concentration mostly occurs in relatively small markets (in terms of gross output), while for the global buckets, the decrease occurs in relatively big sectors.

Finally, with respect to entrenchment, which captures the persistence of firms at the top, the trends remained relatively flat across all geographical buckets in the period considered (2000-2019).

2.2 Literature Review

There is a growing literature documenting recent trends on market concentration, markups, entrenchment, and other measures that can proxy trends in competition in OECD countries. Studies on concentration differ in their definition of a market – both in terms of the product or industry scope and the geographical coverage – and in their specification of firm boundaries across markets and countries, such as through the role of business groups. Studies estimating trends in average markups differ in the granularity of data used (based on firm-level data or estimated at the sectoral level) and in the estimation technique. In this section, the different strands of the literature are reviewed, focusing on studies that document trends in competition and explanations for these trends and their consequences.

2.2.1 Concentration

The literature is characterised by the presence of several indicators to measure concentration. One of the differences depends on whether concentration should be calculated at the sector or “market” level, national or local level, etc. Indeed, an important question still open when looking at concentration trends is how to compute and measure concentration. Different definitions and proxies have led to different conclusions in the literature, and a consensus has not been reached.

This sub-section begins with studies that document concentration trends at the national level, most of which describe an increase in the average (industry level) concentration over recent decades. Then, it discusses recent innovations in the measurement of concentration and how they affect the corresponding trends. Several analyses that measure concentration ratios at the national industry level, mainly focused on the US, have documented an increase in national average concentration over the past few decades (Autor, Dorn, et al. (2020); Barkai (2020); Covarrubias, Gutierrez, and Philippon (2019); Furman and Orszag (2018)). Grullon, Larkin, and Michaely (2019) show that more than 75% of US industries have experienced an increase in industry concentration since the late 1990s. Looking over the past 100 years, Kwon, Ma, and Zimmermann (2023) document a long-term rise in concentration.

In European countries, studies have also found increasing concentration, albeit usually at a slower rate than in the US. For example, Lashkari, Bauer, and Boussard (2019) and De Ridder (2023) both find rising concentration using administrative data for France, and De Loecker, Eeckhout, and Mongey (2022) show similar patterns for the UK. Koltay, Lorincz, and Valletti (2022) look at the five biggest European countries (France, Italy, Germany, Spain and UK) using commercial data (Euromonitor) and document a rise in industry concentration and in the share of industries defined as highly concentrated. In addition, they also show that high levels of concentration are associated with antitrust interventions by the EU Commission. The IMF (2019), using a similar methodology on commercial data (Orbis), documents rising average concentration across 27 economies, both advanced and developing. Affeldt et al. (2021) use a different approach and construct market shares starting from information available from EU Commission merger cases and also find that concentration has increased over time. The comparison between the results obtained with this novel dataset and other results that rely on more typical measures suggests that the level of concentration may be magnitudes higher than that implied by previous studies².

Previous studies on Europe – even those that look jointly at several countries – compute concentration measures at the national level and then study cross-countries dynamics. Bajgar, Berlingieri, et al. (2023) is an exception as it considers the E.U. as a single market and computes concentration measures at the national but also directly at the European level. Accounting for the cross-country subsidiaries of business groups, they document a slight increase in concentration across 12 European countries included in their sample, both at the national and European levels.

In contrast with previous evidence, Gutierrez and Philippon (2023) and Kalemli-Ozcan et al. (2023) have found concentration trends in Europe to be flat or decreasing. These differences are likely to be explained by contrasting methodologies. Importantly, Gutierrez and Philippon (2023) construct the total market size (the denominator of concentration measure) using data from Orbis, which has increasing coverage of small firms over time and can consequently lead to flat industry concentration trends. Kalemli-Ozcan et al. (2023) do not account for connections between firms within a unified business group, which leads to contrasting trends in concentration (see Bajgar, Berlingieri, et al. (2019) for a detailed discussion about these two points)³. As discussed in Section 4, accounting for such connections between firms is important, given the potential implications for the link between concentration and market power. Recent papers on industry concentration have sought to address issues surrounding the appropriate definition of a market and show that,

²One potential concern is that the nature of the data: given that the definition of a market is based on EU antitrust cases, the focus is on problematic sectors and, therefore, those with higher levels and increases in concentration. Furthermore, in 43% of their observations, the Commission reports only the joint market shares of the firms involved in a prospective merger.

³They find increasing concentration trends only when restricting the sample to firms reporting consolidated account. Note that, as carefully explained in Section 4 and in Bajgar, Berlingieri, et al. (2019), the methodology followed in this paper does not use consolidated accounts and relies on unconsolidated accounts.

once the role of international trade is accounted for, concentration in manufacturing remains flat or decreases.

Amiti and Heise (2021) note that existing studies had measured concentration using only sales of firms located in the relevant country rather than sales on the relevant market, i.e., without taking into consideration international trade, in particular import competition. For this purpose, they merge confidential information from three different sources: the U.S. Census of Manufacturing, time-consistent establishment level information at the 6-digit North American Industrial Classification (NAICS) 2007 industry code, and transaction-level import data from the Longitudinal Firm Trade Transactions Database (LFTTD). These data allow them to cover the universe of U.S. imports since 1992, construct the market shares of the foreign sellers in the U.S., and correct for double counting of imports from U.S. plants abroad. Accounting for this import competition, they show that U.S. industry concentration has been flat between 1992 and 2012 because foreign firms have increased their exports to the U.S., even if their individual market shares tend to be small (so, on average, it increases the overall size of the market more than it affects the contribution of the top firms). Concentration mainly fell in industries that had high initial import penetration, which are also the industries that experienced the fastest growth in import competition. In the U.K., the Competition and Authority (2022) (henceforth, CMA) show that correcting concentration ratios for international trade causes a fall in the level of concentration. However, there is still a slight increase over the period 1997-2018⁴.

A key aspect of measuring concentration is to accurately define the geographical scope of a market. For non-tradable products, this is likely to be local geographies, while for tradable products, the relevant market may be international. Here, the evidence is mixed. A growing literature documents that the geographical dimension at which markets are defined affects the implied concentration trends. Rossi-Hansberg, Sarte, and Trachter (2021) observe that the national trend of increasing concentration is not reflected in average local market concentration, which is declining in the US. For products that can only be supplied locally, they argue that the relevant market is sub-national. They explain the differing trends at the national and local levels by observing that large firms are expanding by opening establishments in new local markets. Relatedly, Hsieh and Rossi-Hansberg (2023) document how the “industrial revolution in services” – the increasing returns to fixed-cost-intensive technologies and changing management practices in services sectors – has led to the expansion into new markets and a reduction in local concentration. Rinz (2022) finds similarly decreasing trends in local concentration between 1976 and 2015. However, Autor, Patterson, and Reenen (2023), still looking at the US, find instead that only local employment concentration has decreased, while local sales concentration has increased. They explain the divergence in local and national employment concentration trends with the structural shift of the economy, with a reallocation of economic activity from relatively concentrated manufacturing sectors, where employment concentration is high, to relatively unconcentrated services sectors, where employment concentration is lower. Looking within industry-by-country cells, concentration has increased, even for employment. They suggest that the differences between their findings and those of Hsieh and Rossi-Hansberg (2023) and Rinz (2022) are explained by the use of alternative datasets. Benmelech, Bergman, and Kim (2022) and Smith and Ocampo (2022) find increasing average local concentration, in line with Autor, Patterson, and Reenen (2023), with the latter focusing on the retail sector.

Some studies have argued that the relevant market is supra-national. Lyons, Matraves, and Mofatt (2001) estimate a model of industry concentration that endogenously allows for markets to

⁴Freund and Sidhu (2017) find that an increase in the number of emerging market firms in an industry is associated with a decline in concentration, looking at both manufacturing and services industries.

be defined at either the national or EU level, showing that the four countries studied – France, Germany, Italy, and the UK - varied in their integration with the EU. Affeldt et al. (2021) use market definitions from EU horizontal merger cases to define the geographic scope, showing that concentration increased most in worldwide markets. Note, though, that their sample is an unbalanced panel, and the relevant geography is not fixed over time, so the results could be driven by changing sample composition if more concentrated sectors become global over time.

Most of the literature on concentration has used industries to proxy for markets, mainly due to data availability at this level of aggregation. Industry classifications are constructed to reflect production processes and may not reflect consumer product markets. Industries are often also much broader than product markets. Therefore, some authors have argued that industry concentration may not capture an economically relevant measure of market concentration (Berry, Gaynor, and Morton (2019); Werden and Froeb (2018); Shapiro (2018[28]). Benkard, Yurukoglu, and Zhang (2021) re-examine trends in US concentration using data that more accurately reflect consumer product markets, although the set of markets covered is more limited. They find that the level of concentration is typically higher in product markets, with 45% of the industries in their sample defined as highly concentrated according to Horizontal Merger Guidelines. However, they show that the median product market concentration decreased between 1994 and 2019. In line with previous studies, when they aggregate to the industry level, concentration increases. They explain these divergent trends with a model in which the cost of firms for supplying in adjacent product markets has fallen over time, so firms expand by adding products in new markets within the same industry. The model of Aghion et al. (2019) follows similar lines and may also be able to explain these trends. Pellegrino (2023) takes an alternative approach to defining relevant product markets based on observed similarities and substitutabilities across products, with the results suggesting a broad increase in market power over time.

2.2.2 Entrenchment

Market concentration is a static measure, as it captures the market shares of leading firms but does not identify whether there is churn among the market leaders. Schumpeterian “creative destruction” and industry dynamism among market leaders can also reflect competition (for the market), even in markets characterized by high concentration levels. Therefore, an important measure that may reflect a lack of competition is the extent of entrenchment among market leaders.

Some studies have used dynamic measures of firms’ entrenchment in markets. Bessen et al. (2020) proposes two alternative measures of entrenchment: first, the annual displacement hazard rate of firms ranked in the top 4 in an industry dropping out of the top 4 (i.e. the likelihood that a leading firm loses its place among the top firms in a market); second, the annual hazard rate of a firm ranked fifth to eighth in the industry progressing into the top four (i.e. the probability that a competing firm leapfrogs a top 4 firm to become a market leader). Using US data, he finds that displacement rates of the top firms rose from the 1970s to the 2000s but have declined sharply since then. Freund and Sidhu (2017) use Orbis data to show that between one-third and one-half of firms that were in the top four in an industry in 2014 are different from those in 2006. Competition and Authority (2022) finds that, in the UK, the likelihood that the ten largest firms in an industry were also among the largest ten firms three years before has increased over the last two decades, implying that competition among market leaders may have fallen. Davies (2022) shows an increasing persistence of the largest firms in the top 10 until 2018 in the UK. Furthermore, the study finds that the persistence of the largest firms at the top is more pronounced in more concentrated industries. On the contrary, using data from Japan, Honjo, Doi, and Kudo (2018) show that market leaders are more likely to be replaced by competitors in industries with negative

growth and high concentration. Bajgar, Criscuolo, and Timmis (2021) focused on three variables to explore the churning of the top firms: the share of firms that are in the top 8 but were not in the top 8 in the previous year, the rank correlation between the market shares of top 8 firms over two years, and the market share instability. Using data at the country and industry level for a sample of OECD countries, they show that increased concentration is associated with reduced churning among the top firms, namely with less entry of new firms at the top and more ranking persistence of the leading firms.

2.2.3 Explanations of recent trends in competition proxies

Numerous theoretical and empirical explanations have been proposed to explain the observed increasing trends in industry concentration and the entrenchment of industry leaders. Many claim that these trends are, at least partially, driven by technological change, including the rise of intangible capital and lower diffusion of technology between firms, while others argue for institutional factors such as antitrust policies or declining worker power. Many, or all, of these explanations likely play a role. Some authors have also discussed the extent to which these macro trends do indeed reflect increasing market power; Syverson (2019) suggests that market power is a leading candidate explanation of these trends, while Berry, Gaynor, and Morton (2019) highlight the importance of establishing causal relationships to understand the drivers of these trends.

The fall in IT costs could disproportionately affect market-leading firms. Aghion et al. (2019) propose a model in which the cost of supplying to multiple markets has fallen due to IT advances, causing the most efficient firms to expand into new markets. Even though markups fall within firms, the reallocation of activity to high-markup firms causes an aggregate increase in markups. Using data on French firms, Lashkari, Bauer, and Boussard (2019) document a positive within-industry relationship between firm size and IT demand, and this disproportionately affects larger firms through mitigating issues of organisational efficiency associated with firm growth.

More generally, intangible capital, which incorporates IT technologies, can explain trends in market power. De Ridder (2023) models intangibles as a fall in marginal costs and a rise in fixed costs. Intangible-intensive firms can operate at low marginal cost and, therefore, deter entry into their market. These firms also have greater incentives to innovate, causing an increase in overall R&D. However, its benefits are lower because it is concentrated in a few leading firms. Crouzet and Eberly (2019) also focus on intangibles, noting that the combination of their scalability and legal protections (i.e., patents) can lead to a rise in concentration. Studies show that intangibles are correlated with concentration, markups, business dynamism and entrenchment (Bajgar, Criscuolo, and Timmis (2021); Calvino, Criscuolo, and Verlhac (2020), Calligaris, Criscuolo, and Marcolin (2018), Berlingieri et al. (2020), Bessen et al. (2020), Covarrubias, Gutierrez, and Philippon (2019)). On the other hand, De Loecker, Eeckhout, and Unger (2020) discuss the possibility of the reverse relationship, whereby an increase in concentration can cause a decrease in investment in intangibles because a lack of competition can reduce the incentives to innovate and invest.

A slowdown in technology diffusion from leading innovative firms to followers can also explain the macro trends on concentration and business dynamism. Akcigit and Ates (2021) propose such a model, in which diffusion has reduced because firms have increasingly used intellectual property rights to deter technology transmission outside the firm. Olmstead-Rumsey (2020[53]) shows that the decline in innovation of small firms can explain the rise in concentration, providing evidence that small firms' patents have made less significant innovations in the 2000s relative to the 1990s. Andrews, Criscuolo, and Peter N. Gal (2019) provide evidence of increasing productivity divergence between firms at the global frontier of productivity and laggard (non-frontier) firms. This divergence can potentially be triggered by structural changes in the global economy, such as digi-

talisation, globalisation and the rising importance of tacit knowledge, that fuel rapid productivity gains at the global frontier.

Autor, Dorn, et al. (2020) propose that technological change, combined with globalisation, can disproportionately benefit the most efficient firms, leading to a rise in concentration and higher markups (the latter because most efficient firms have higher markups on average). Kwon, Ma, and Zimmermann (2023) suggest a positive relationship between rising industry concentration and technological intensity and higher fixed costs, as well as suggesting that globalisation has played a role in rising concentration in recent decades. Antoniadou (2015) develops a model in which an increase in market toughness, such as due to globalisation, causes an increase in competition but also an increase in the scope for quality differentiation, which generates an incentive to invest in fixed costs of innovation, leading to increased quality, markups, and prices.

In contrast, more relaxed antitrust policies and merger enforcement could provide an alternative explanation for increasing competition and entrenchment. In the US, Grullon, Larkin, and Michaely (2019) show that industries that experienced the largest increases in product market concentration had more profitable M&A deals but no increase in operational efficiency. Gutierrez and Philippon (2023) and Döttling, Gallardo, and Philippon (2017) argue that there has been a divergence in the strength of antitrust enforcement between the US and EU. In Europe, antitrust enforcement has been stricter, and product market regulations have decreased, in contrast to the US. These differences explain the higher investment and declining real prices in the EU⁵.

2.2.4 Contribution of this report to the literature

This project makes numerous contributions to the literature on competition in OECD countries. It covers in an integrated way multiple measures that proxy for different facets of competition (or a lack of it): concentration and entrenchment. Examining different characteristics of markets in the same setting provides richer inferences about the trends in market power, as well as possible explanations for these trends. The analysis is conducted for both manufacturing and services (and beyond, as it also includes mining and utilities sectors) across many OECD countries and over a long-time horizon. Constructing a database that allows such analysis is itself a significant contribution.

The project innovates on the existing literature in its measurement of concentration, aimed at reflecting markets more accurately. First, it measures concentration within more narrowly defined industries than most previous cross-country studies, mainly at the 3-digit level. Second, international trade is accounted for by constructing a taxonomy that defines whether markets compete domestically or internationally and computes concentration at the corresponding geographic level. Using this taxonomy means that the concentration measures account for firms' international activities, even when firm-level trade data is absent. On top of this, in the robustness checks, imports and exports are incorporated in the concentration measure using industry-level data. Third, following the methodology by Bajgar, Berlingieri, et al. (2019), the connectedness of firms within business groups is accounted for to incorporate the role of mergers and acquisitions in driving concentration trends and to capture the complete activities of multinational firms in a market.

⁵Various alternative explanations for changing indicators of competition have also been put forward. A reduction of worker power can lead to a redistribution of rents from workers to firms, causing a rise in corporate profitability and a fall in the labour share (Stansbury and Summers (2020)). A decline in long-term interest rates can disproportionately increase investment by market leaders relative to followers (Liu, Mian, and Sufi (2022)). A decline in population growth (Peters and Walsh (2020)), growth of industries (Kwon, Ma, and Zimmermann (2023)), economic growth (Ekerdt et al. (2023)), or network effects (Berry, Gaynor, and Morton (2019)) have all been proposed as alternative explanations.

Alongside concentration, a static measure of market shares at any point in time, the project also measures entrenchment. Entrenchment is a dynamic measure of the persistence of firms as market leaders and provides richer insights into the extent of competition, even when concentration is high. The measurement of entrenchment innovates on previous literature by defining markets tradable internationally following the taxonomy and by accounting for the connectedness of firms within a business group.

2.3 Data

2.3.1 Overview

The analysis requires merging detailed data collected from numerous sources. The data needed include product level and firm-level data, as well as data at the industry-country-year level, with industries generally defined at the 3-digit NACE revision 2 level of granularity⁶. Using cross-country data at this level of granularity combined with international trade data is a key innovation of this study.

The total production value for each industry-country-year is a key variable used in constructing both the concentration measures and the geographic taxonomy of industries. Production data is collected from National Accounts (NA), the STAN database and Eurostat's Structural Business Statistics (SBS).

Production data are matched with data on international trade flows from OECD Inter-Country Input-Output (ICIO) tables, the "Base pour l'Analyse du Commerce International" (BACI) dataset, and the Trade in Services by Partner (TISP) data. The importance of trade flows is twofold: i) they are necessary to define a novel taxonomy developed to identify the geographic level at which industries compete, and ii) to carry out relevant robustness exercises of concentration measures that account for trade flows. Another novel contribution of this study is to use trade data in a cross-country setting to construct such a taxonomy and adjust concentration measures.

Firm-level production data collected from Moody's Orbis database is required to measure the largest firms' contribution to total production and to establish their relative position in each market-industry ranking. The analysis relates proxies of competition to various economic variables, such as foreign ownership, intangibles, and policies, which are collected from numerous sources described below.

For concentration measures and entrenchment proxies, the final sample covers 15 European countries plus three non-EU countries (Japan, South Korea and USA)⁷. Regarding sectoral coverage, the analysis focuses on mining, manufacturing, non-financial market services (excluding real estate) and utilities sectors following the NACE Rev. 2 classification at the 3-digit industry level. Due to data limitations and different data requirements to construct each alternative proxy of competition analysed in the study, the sector level of aggregation and coverage slightly changes across different sections of the study. Some industries have been aggregated at a higher level with respect to the full list of 204 3-digit industries, and some industries have been excluded from the analysis due to data limitations, leaving a baseline sample of 127 industries. Out of these, 112 (88%) are 3-digit, 10 (8%) are 2-digit, and 5 (4%) are an aggregation of two or more 2-digit industries⁸. The period covered is 2000-2019.

This section discusses the datasets used in the analysis, including the processes conducted to clean and prepare the data.

⁶NACE is the "statistical classification of economic activities in the European Community" and is the subject of legislation at the EU level, which imposes the use of the classification uniformly within all the Member States. The present NACE Rev. 2 is the new revised version of the NACE Rev. 1 and of its minor update NACE Rev. 1.1. In this study, industries are generally defined at the 3-digit level of aggregation; however, in some cases, industries must be aggregated due to data constraints, as described in this section.

⁷The countries with suitable data quality are: Belgium, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Norway, Poland, Portugal, Slovenia, Spain, Sweden, and United Kingdom

⁸For further details on the different levels of aggregation in the concentration sample, please refer to section A of the Appendix

2.3.2 Production data

The total value of production, defined as gross output in millions of euros in each industry-country-year, is a key variable used for constructing the concentration measures and entrenchment proxies, as well as the geographic taxonomy of industries. In order to get proxies of industry concentration measures and entrenchment granular enough to be as close as possible to market level measures, a substantial effort has been made in this study to calculate the value of production at the 3-digit NACE Rev. 2 level.

The main data source for gross output is the Eurostat National Accounts (NA, henceforth), which is the primary dataset used by countries to measure Gross Domestic Product (GDP) and other key economic variables⁹. However, Eurostat publishes NA, including the value of production, at the A*64 level of the industry classification NACE Rev.2, comprising of aggregations of 2-digit level activities¹⁰. It is therefore supplemented with data from Eurostat's Structural Business Statistics (SBS), which contain information on the economic activity of all economic sectors excluding agriculture and personal services, and provide data on the value of production at the 3-digit level of aggregation¹¹.

In order to get production data at the 3-digit level, the 3-digit SBS data are used to construct the share that each 3-digit industry represents within its own 2-digit industry. These shares are then used as weights to apportion each 2-digit production value from the NA to the corresponding 3-digit industries. The 3-digit production data obtained are therefore consistent with NA – which are often based on the population of firms – and at the same time available at the desired level of granularity. The main reason for using the SBS 3-digit production data to apportion the 2-digit ones coming from NA instead of relying directly on them is the following: SBS captures the structure of the economy at a higher level of disaggregation than NA, and data are representative of the economy within countries but not across countries (due to different methodologies adopted by National Statistical agencies to collect them). Therefore, SBS is an excellent source for obtaining long time series of 3-digit weights, but less ideal to directly measure production of the population of firms in a cross-country context. Thus, to guarantee cross-country comparability, NA are preferable when looking at totals. There are two main obstacles to overcome for obtaining consistent 3-digit level data on production over the long period considered. First, the SBS dataset contains missing values (about 26% and 9% in the pre-2008 and post-2008 samples, respectively), so an imputation procedure is needed. Second, the classification of economic activity changed in 2008 (from NACE Rev. 1.1 to NACE Rev. 2), and SBS provides two different time series: one for the years pre-2008, reported in NACE Rev. 1.1, and one for the years post-2008, reported in NACE Rev. 2 (this one available from 2006). Therefore, a conversion from the old to the new NACE classification system is required to have a time series at the industry NACE Rev. 2 level from 2000 to 2019.

Where possible, a multi-step imputation procedure has been set up to fill in missing values in the SBS data. The various steps are imposed sequentially, according to the strength of assumptions required, with priority given to the most robust methods. The next paragraph provides a brief description of each method (and parentheses report the weighted average percentage -across the two SBS samples -of observations recovered through each of them)¹². For further details, both on

⁹For the extended documentation see: National Accounts metadata; for further definitions see: European system of accounts - ESA 2010

¹⁰NACE Rev.2 is fully compatible with ISIC Rev.4. Please note that in this section 2-digit will be used also to indicate generic A*64 industries that are sometimes slightly more aggregated. For further details on the A*64 classification, see OECD industry classification

¹¹For further details please refer to Methodological manual on European Structural Business Statistics – 2021 edition and the SBS website Structural business statistics

¹²The weights are given by the relative size of the two samples (pre-2008 41%, and post-2008 59%) in the total


the methodology and on the imputation process results, see Annex A.

First, accounting identities are used, whereby if only one 3-digit industry value is missing within a 2-digit industry, the missing value is filled with the difference between the 2-digit and the sum of the non-missing 3-digit (2.2%). Second, interpolation is used if the previous year and following year values are available (1.4%). Third, a method similar to propensity score matching is implemented, which uses growth rates in similar countries² to impute missing values (3.2%). In this context, for a given industry, two countries are considered similar if they grew approximately at the same pace in the previous year. Then the average growth of the five closest countries (using the aforementioned growth criterion) is used to impute the missing value. Finally, regressions using additional non-missing economic variables, such as turnover, employment, and the number of firms, are used to predict missing values (3.9%). If, following these steps, a value is still missing, it is left as missing.

The NACE classification was updated from NACE Rev. 1.1 to NACE Rev. 2 in 2008. Eurostat provided correspondence tables to facilitate conversion between the two coding systems. However, the changes were substantial, and the correspondence between the two classification systems is many-to-many (i.e., there are multiple correspondences between industries). Given that the present analysis requires a consistent and unique industry classification -chosen to be NACE Rev. 2 -data pre-2008 had to be converted to the newer system. Even in the most obvious cases, such as industries with one-to-one mapping from Rev. 1.1 to Rev. 2, the values in the overlapping year in the two datasets, 2008, are not always consistent. Therefore, to obtain the full time series, the pre-2008 data are used to compute country-industry growth rates across consecutive years, which are then applied backwards using the 2008 values obtained from NACE Rev. 2 as a starting point. Further details of the conversion from NACE Rev. 1.1 to NACE Rev. 2 are reported in Annex A.

As mentioned above, SBS data are used as weights to apportion the values provided by NA at the 2-digit industry level into 3-digit industries, while ensuring that that the values of all the 3-digit within a 2-digit As mentioned above, SBS data are used as weights to apportion the values provided by NA at the 2-digit industry level into 3-digit industries, while ensuring that that the values of all the 3-digit within a 2-digit industry still aggregate to the NA figures. The apportioning procedure is diagrammatically represented in Figure 2.1

Figure 2.1: Apportioning for Production Data



Note: the chart illustrates the process of apportioning aggregated production data into more disaggregated industries.

SBS sample.

Mathematically, let $GO_{sct}^{2d NA}$ be the value of gross output for the 2-digit industry S in country c and year t obtained from the NA data. Then let GO_{sct}^{3dSBS} be the value of gross output for the 3-digit industry S in country C and year T obtained from SBS. The apportioned values of gross output at the 3-digit level GO_{sct}^{3dNA} are calculated as:

$$GO_{sct}^{3d NA} = \frac{GO_{sct}^{3d SBS}}{GO_{Sct}^{2d SBS}} GO_{sct}^{2d NA}, \quad (2.1)$$

where the first term of the right-hand side are the weights calculated from SBS data, with $s \in S$ representing each 3-digit industry contained in the associated 2-digit industry S , such that $\sum_{s \in S} GO_{sct}^{3dSBS} = GO_{Sct}^{2dSBS}$. The computed value GO_{sct}^{3dNA} is the measure of gross output used throughout the analysis.

The data for extra-EU countries included in the analysis (Japan, Korea, and the United States) is required for industries that compete at the global level, as defined by the taxonomy. In most cases, these data are available at the desired level of disaggregation in the original data sources used and, as such, no additional cleaning or data preparation was needed¹³. However, for two industries in the services sector, data was only available at the 2-digit level, with no 3-digit level production data attainable from alternative sources. In these cases, exports were used as a proxy of the shares of economic activity attributed to each 3-digit industry within its associated 2-digit industry (equivalent to the procedure described above, except the more detailed data used was exports). The assumption that production shares are proportional to export shares is stronger, as some industries may be more tradeable than others, but exports are positively related to production, so this provides a reasonable approximation.

Annex A provides detailed information on the NA and SBS datasets, the imputation procedure, the mapping of pre- and post-2008 data, as well as the apportioning procedure. Finally, in Table 2.10, there is the list of the 151 sectors included in the final sample of production and of the 127 included in the analysis of concentration, leadership ratio, and entrenchment.

2.3.3 Trade data

Trade data at the 3-digit level (or slightly more aggregated due to data limitations) are needed to define the taxonomy that identifies the geographic level at which an industry competes. Furthermore, as a robustness check, the denominator in the concentration measures is adjusted to consider import penetration and exports.

To obtain data on import and export flows at the 3-digit industry level, international trade data is collected from three main sources.

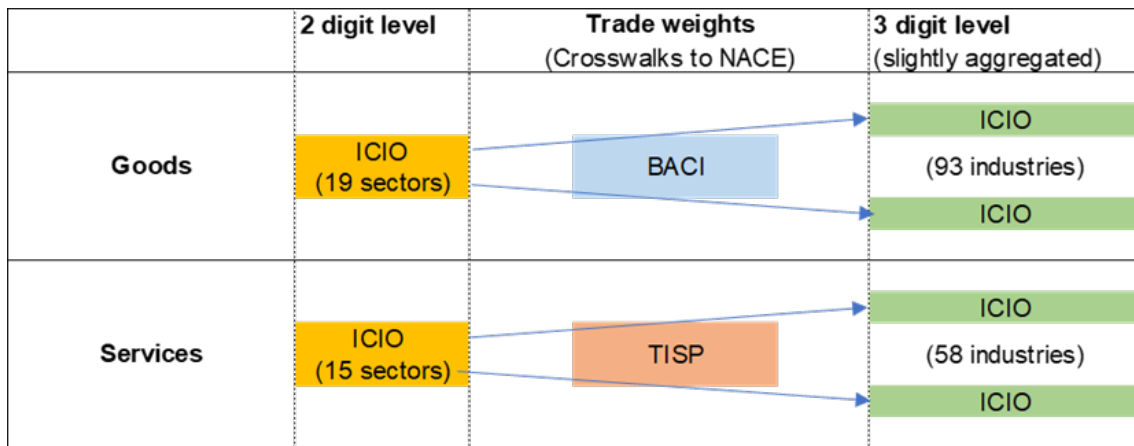
The primary data source is the OECD “Inter-Country Input-Output” (ICIO) tables, which provide import and export flows for each country, industry, and year, and can distinguish trade with EU and non-EU countries. The ICIO data have two main advantages over other standard trade data. First, they account for re-export, meaning that trade is not disproportionately allocated to countries with large ports for redirecting shipments, such as Rotterdam in the Netherlands, for instance. Second, they are in basic prices, in line with the data on production used in the analysis, meaning that comparisons can be made more accurately. However, ICIO data are only available at a higher level of aggregation (more aggregated than the 2-digit level) than the desired one.

¹³Data for the extra-EU countries come from the OECD STAN database complemented with data from the Korean Statistical Information Service (Korea), from the Ministry of Economy, Trade and Industry (Japan), and from the Bureau of Economic Analysis (United States). Period covered: 2000-2019.

These data are thus supplemented with highly disaggregated information on trade flows from the Centre d'études prospectives et d'informations internationales (CEPII) "Base pour l'Analyse du Commerce International" (BACI) database for physical goods, and the OECD "Trade in services by partner economy" (TISP) database for non-financial market services. These data are reported at the product level; therefore, crosswalks have been used to convert the product level data to disaggregated industry level. Following this matching exercise, it has been possible to recover goods trade data for all manufacturing and mining industries at 3-digit level, while industries belonging to utilities and non-financial market services have, in some cases, been aggregated at high level¹⁴. However, they do not have the advantages of being in basic prices and accounting for re-export, so they are not used directly.

The detailed trade data are used to construct weights at the country-partner-industry-year level, where the partner is usually either EU or non-EU countries, and an apportioning procedure in the same fashion as that used for the production data is carried out for the trade data. The weights are used to distribute the more aggregated trade values from ICIO to more detailed industries. This method maintains the desirable characteristics of the ICIO trade flows data whilst providing a more disaggregated industry classification. The process is represented diagrammatically in Figure 2.2¹⁵.

Figure 2.2: Trade data apportioning



Note: the chart illustrates the process of apportioning aggregated trade data to disaggregated industries.

The apportioning procedure can be described more formally as follows. $Trade_{Sct}^{2d, ICIO}$ is observed in the ICIO data, where S denotes the 2-digit (or more aggregated) industry, c denotes the country, and t is the year. $Trade_{Sct}$ can be imports or exports, and with either the EU or non-EU as a partner, for a 3-digit industry s .

¹⁴Out of the 99 3-digit industries belonging to utilities and non-financial market services, it has been possible to recover trade data at 3-digit level for 37 industries, while the remaining 62 have been aggregated into 14 2-digit industries (or slightly more aggregated levels). For the full list of industries used, see Table 2.10.

¹⁵Note that with this method, it is possible to have exports larger than production for some observations even if at the 2-digit level it is imposed by the production and ICIO data that exports values are lower than total production. This can happen when the 3-digit trade share is high relative to the 3-digit production share, and the production and trade values at the 2-digit level are relatively close. This is the result of noise in the data. For example, for multi-product firms, production is typically allocated to their main (single) industry by statistical offices, while their trade flows would be split across products in customs declarations or in services trade surveys. The more disaggregated the industry definition, the more likely discrepancies can arise.

Let $s \in S$ be any 3-digit industry s included in the 2-digit industry S . From BACI and TISP, $Trade_{Sct}^{3d}$ is observed, where s are more disaggregated industries that can be aggregated to the 2-digit level S . Then trade at the 3-digit industry s level is defined as:

$$Trade_{sct}^{3d ICIO} = \frac{Trade_{sct}^{3d}}{Trade_{sSct}^{2d}} Trade_{Sct}^{2d ICIO} \quad (2.2)$$

Annex A provides further details about the international trade datasets and the cleaning procedures applied in each case¹⁶.

2.3.4 Firm-level financial data

Financial data at the firm level are required to measure both concentration and entrenchment. The primary firm-level dataset used is the Moody's Orbis database. It collects data on 450 million listed and unlisted firms worldwide, providing information on their economic activities, such as employment and output. The data used in the analysis are obtained from the 2021 Orbis vintage. Numerous steps are undertaken to clean the data, closely following Peter N Gal (2013) and Kalemli-Ozcan et al. (2023).

Financial information within Orbis is available at the business-group level (consolidated financial data aggregated across all firm subsidiaries) and at the individual firm level (unconsolidated information referring to individual firms' activities). In this report, as motivated in Section 4, unconsolidated accounts are used¹⁷.

For the analysis on concentration and entrenchment, further data are collected from WorldScope, provided by Thomson Reuters, a cross-country firm-level commercial database of listed firms. This dataset covers 95% of the global stock market capitalisation and allows to substantially increase the coverage of listed firms included in the sample. These additional data provide relevant information, especially for non-EU countries, usually less represented in Orbis¹⁸.

The Data Appendix provides a detailed summary of the steps taken to clean and prepare the Orbis data for the analysis.

¹⁶Note that the trade data capture trade in the final industrial activity of each industry at the point at which the good or service crosses the border. As an example, considering the car industry, the trade variables capture the value of cars crossing the border from the source to the destination country. The trade variables do not capture trade in all products of firms classified in a particular industry. For example, a firm in the car industry may import tires, windscreens, and other parts – these would be classified under the other relevant 3-digit level industries. Furthermore, only products that cross borders are captured in the data: no sales of a car manufacturer located in the UK that only sells to the UK market would be registered as trade in cars in the data.

¹⁷As explained carefully in Section 4, the measures of concentration and entrenchment built in this report look at business group activities. The approach adopted fundamentally relies on unconsolidated data of the individual subsidiaries within a business group, as the objective is to identify the precise industry and location of all the subsidiaries belonging to a group and to correctly apportion the group sales to the markets in which the business group is active. In Annex A, additional details on the data cleaning and preparation are provided, especially about when and how consolidated accounts are used.

¹⁸As explained in Bajgar, Berlingieri, et al. (2019), Orbis data are generally well suited to analyse industry concentration in Europe – the main geographic market of interest in this study – since it has a good coverage of medium and large enterprises in these countries. In contrast, Orbis generally has poorer coverage of subsidiary-level information for non-EU countries, especially for US firms. Note that this concern does not apply to business group-level information (consolidated accounts), since Orbis and WorldScope together cover close to the universe of listed firms.

2.3.5 Firm ownership information

For the analysis on concentration and entrenchment, the key variable of interest is the gross output of business groups – rather than firms – in each industry (for industries in the European and global buckets) or industry-country (for industries in the domestic bucket), following the methodology developed by Bajgar, Berlingieri, et al. (2023) (see Section 4 for a detailed discussion on the choice of looking at business groups). Indeed, it is necessary to reconstruct the worldwide structure of the business group, identifying ownership linkages between headquarters and all their subsidiaries. Orbis data contain detailed ownership information, with each firm being linked to its global ultimate owner. The global ultimate owner is defined as the firm owning at least 50.01% of the total shares of a subsidiary. However, Orbis ownership information only covers the period 2007-2020. To get this information for years before 2007 and to further complement and quality-check the existing information, Orbis is supplemented with data from the Zephyr database, also provided by Moody. This database reports information on the Merger and Acquisition (M&A) activities of firms around the world. It captures both domestic and cross-border M&As and covers deals involving target and acquiring firms across all industries. Observing M&A activity between firms enables to track changes in ownership over the years, allowing to identify the business group structure also before 2007. A detailed discussion of the methodology is provided in Annex A, together with further details on the process of cleaning and harmonizing multiple data sources. For a very detailed explanation of the methodology used to build the business group structure, refer to Bajgar, Berlingieri, et al. (2023)

2.4 Methodology

2.4.1 Overview

This section describes the methodology used to define the competition proxies considered in the report: concentration measures and entrenchment. Each of these proxies provides some information about the competitive environment of the (broadly defined) market in which they are computed. Analysed together, they provide some indication of the extent of market power held by the leading firms and of competition in each market.

Market concentration is the main focus of competition authorities. However, when measuring concentration, a major empirical issue is how to get as close as possible to the definition of the “relevant market”. First, defining the relevant market can be conceptually difficult. Second, data limitations arise when trying to measure it. As a result, industry concentration measures are often used as a proxy for market concentration.

However, industry concentration is related to but distinct from the concept of market concentration. The fact that a large share of industry activity is due to a handful of leading firms does not necessarily mean that product markets within an industry are highly concentrated (OECD (2021)). Market concentration is a narrower definition than what it is typically reflected in industry concentration measures. Despite the important limitations of industry concentration measures, they continue to be widely used, data limitation being the major obstacle preventing going to a narrower level of detail. Throughout the baseline analysis, markets are defined using industry classifications. One important innovation of the project is to use more narrowly defined industries than previous cross-country studies, going from the 2-digit to mostly 3-digit level. Moreover, in section 6, the report shows that at the level of aggregation and for the countries considered in the report, concentration at the industry level is correlated with product-level concentration, and they exhibit similar trends on aggregate. Therefore, in this context, concentration measures at the industry level are likely to also indicate consumer product market concentration.

The world economy is also globalised, with markets integrated across countries. The value of international trade amounts to more than half of global GDP every year, and the stock value of outward investments is 44% of world GDP¹⁹. “Global firms” operate in many countries, either by trading internationally or establishing affiliates (see, for example, Bernard et al. (2018)). This globalisation of firms also affects the appropriate market definition from a geographical perspective. This project innovates with respect to the existing literature by accounting for the international dimension of competition when defining markets, as well as the ownership and realm of operation of firms. It uses cross-country data on the activities of firms, the size of industries, and international trade in each industry to define the geographical level at which markets compete. Specifically, it constructs a taxonomy of sectors that defines whether each industry competes at the domestic, EU, or international level. This taxonomy is then used to account for international competition in calculating proxies of concentration and entrenchment.

To summarise, in this study, a significant effort has been made to get as close as possible (given data limitations) to a market definition concerning both the industry level of aggregation and the geographic level of competition. First, most of the analysis is conducted at a 3-digit industry level, allowing a more disaggregated level of analysis with respect to previous cross-country studies and a finer market definition. Second, regarding the geographic dimension, each industry defines the relevant geography as either domestic, European or global based on the outcome of the taxonomy.

¹⁹See World Bank national accounts data, Trade (% of GDP) and OECD data, Foreign direct investment (FDI) stocks.

For the remainder of the paper, the term geography will refer to the regional span determined by the taxonomy, and the term market will identify the combination of industry and geography in which firms operate.

A final contribution of the study is to document patterns in multiple competition measures in the same setting. Concentration is evaluated by looking at both the concentration ratio, which captures the share of production accounted for by the four largest firms in a market, and the leadership ratio, which measures the ratio of the sales of the market-leading firm relative to other sets of firms in the market. In addition to concentration, a further measure of competition is computed and analysed, entrenchment, which captures the extent of churn in the leading firms in each market. More precisely, it measures the share of firms that were market leaders in the previous period that remain market leaders in the following period.

This section first describes the methodology to construct the taxonomy, which defines the geographic level at which each industry competes. Then, it describes the definitions of concentration, leadership ration, and entrenchment measures.

2.4.2 Taxonomy of industries

This project extends the existing literature by developing a taxonomy of industries aimed at defining the geographical level at which competition takes place. This is the correct level of geographic aggregation at which industry concentration measures should be calculated to get as close as possible to the definition of “relevant market”.

For industries that are non-tradable, competition takes place mainly domestically²⁰. Consequently, concentration measures should be computed at the national level. For tradable industries, competition takes place mainly internationally, across borders; firms in one country can supply their product to consumers in another. In this case, when looking at concentration trends, it is more relevant to look at the top firms in international markets. At the same time, even in industries that show high volumes of trade, there are different boundaries in the international span that firms can reach due to geographical limitations, trade restrictions, and other factors. Thus, in the taxonomy, tradable industries are defined as competing at either the EU or at the global level, depending on the relevant boundaries of the market.

Data on trade flows are used to identify the geographic dimension at which each industry competes. Trade flows provide a measure of the extent to which firms compete across borders and the availability of foreign products to consumers. Trade flows have been used to determine the geographic level of competition in previous research (Lyons, Matraves, and Moffatt (2001)). The characteristics of each industry determine its trade flows and, therefore, its geographic level of competition: its technological feasibility of supplying or purchasing the good or service across borders, its ability to separate the location of production from that of consumption, as well as trade policy barriers, drive the distinction between domestic, EU, and global markets. The distinction of tradable industries between competing at the EU level and globally arises from the fact that EU countries have low costs of trading with each other relative to trading with the rest of the world. Lowe costs arise from geographical proximity, socio-economic similarity, and the Single Market and Customs Union.

²⁰It could be argued that competition takes place on a sub-national level for some non-tradable services. Unfortunately, sub-national data are not available at the necessary levels of aggregation across the sample of countries. Note that European countries are typically small relative to the US, where the existing literature on sub-national competition has focused (Rossi-Hansberg, Sarte, and Trachter (2021)), reducing the distinction between national and local geographic markets.

The taxonomy is necessarily constant across countries: an industry that competes domestically in one country also competes domestically in other countries, while an industry that competes internationally in one country also competes internationally in other countries. The taxonomy is also time invariant: it is unlikely that a non-tradable industry in one year becomes substantially more tradable in subsequent years and vice-versa, and time-invariance is necessary to ensure that the observed trends in the analysis are not driven by changes in the definition of market²¹. In other words, to meaningfully examine trends in proxies of competition over a long time horizon, the level of aggregation needs to be constant over the period considered. Therefore, the taxonomy is compiled using data aggregated over the full sample period (2000 to 2019).

The taxonomy considers both the imports and exports in each industry from the perspective of firms in the EU to construct measures of openness to trade (OTT, henceforth). Analogous measures of openness to trade are commonly used, especially at the country level (Leamer (1988)). Comparing total trade with domestic production provides an estimate of the tradability of each industry. Specifically, the shares of domestic production, exports, and imports can be decomposed as follows:

$$\text{Domestic share: } \frac{\text{DomesticSales}}{\text{DomesticSales} + (\text{Export} + \text{Imports})_{EU} + (\text{Export} + \text{Imports})_{non-EU}} \quad (2.3)$$

$$\text{EU share: } \frac{(\text{Export} + \text{Imports})_{EU}}{\text{DomesticSales} + (\text{Export} + \text{Imports})_{EU} + (\text{Export} + \text{Imports})_{non-EU}} \quad (2.4)$$

$$\text{non-EU share: } \frac{(\text{Export} + \text{Imports})_{non-EU}}{\text{DomesticSales} + (\text{Export} + \text{Imports})_{EU} + (\text{Export} + \text{Imports})_{non-EU}} \quad (2.5)$$

Note that “domestic sales” is defined as the value of domestically produced output sold to domestic consumers. In the data, it is measured as total production minus exports. Importantly, the three components sum up to 1. Openness to trade is a simple and effective measure of the technological and policy-related feasibility of supplying or purchasing a good or service to a foreign market. Accounting for both imports and exports in a balanced manner provides an estimate of the feasibility of cross-border trade in each industry. Therefore, it provides a relevant metric with which to construct the taxonomy.

An alternative definition of the taxonomy would take a consumption-based approach, whereby exports are not incorporated as consumers are only directly affected by goods and services available in their market, which is captured by domestic production and imports.

The taxonomy is constructed using data from EU countries. It is designed to be representative of the aggregate EU (weighting by the denominator)²². Different time periods (only later years) and weightings (production-weighted and unweighted) are tested for robustness on the taxonomy specification.

²¹Services industries are most likely to have become more tradable over time due to technological advancements and increased integration of trade policy within the Single Market. Therefore, as a robustness check, the taxonomy is constructed on only the later years of the sample (period 2012-2019).

²²Specifically, country-industry-year level data are aggregated across countries and years to get to the industry level, weighting by the denominator of the measure such that those statistics are representative of the EU aggregate.

Taxonomy Thresholds

To assign each industry to a unique category in the taxonomy - i.e., competing domestically, with EU countries or globally - it is necessary to define thresholds on the OTT measure defined above. The methodology implemented for this purpose involves two consecutive steps. First, industries are defined as either competing mainly domestically (non-tradable) or internationally (tradable) depending on the OTT domestic share (note that one minus the domestic share is the traded share). Second, the EU and non-EU shares are compared to assign whether a tradable industry competes mostly at the EU level or globally.

1) Identifying tradable and non-tradable industries The primary objective is to define the threshold that determines whether each industry competes mainly domestically or internationally. The proposed methodology begins with the well-established idea that most manufacturing industries are tradable and, hence, compete internationally. In many conceptualisations, the entire manufacturing sector is assumed to be tradable while services sectors are non-tradable (see, for example, Besley, Fontana, and Limodio (2021), Eaton et al. (2016)). This notion can be used, along with the distribution of domestic shares for the goods sector, to define a threshold for an industry to be considered tradable. Goods sectors are defined as those involving physical products and are covered in manufacturing and mining, while services refer to all other sectors included in the sample (excluding finance and government services). Of course, it is likely that not all manufacturing industries are tradable if, for example, they have very high policy barriers to trade, such as dairy, or high physical trade costs, like concrete. Therefore, the assumption is taken that 90% of industries in manufacturing and raw materials sectors are tradable. Gervais and Jensen (2019) use a similar approach to define whether each industry is tradable, making an assumption that the majority of manufacturing industries should be tradable to identify a threshold for tradability that can be applied to all sectors.

Looking at the distribution of the OTT domestic share for goods industries in the sample, the 90th percentile industry has a domestic share of 0.88 (shown in Table 2.1). Under the assumption that 90% of goods industries are tradable, an industry should have a domestic share below 0.88 for it to be defined as tradable. In other words, an industry is defined as competing mainly domestically if the output sold to domestic consumers is more than 88% of the total output produced, imported, and exported. Therefore, in these cases, international trade is relatively small in the industry. Equivalently, exports plus imports must comprise less than 12% of the sum of the output produced, exported, and imported.

Table 2.1 and Table 2.2 present summary statistics for, respectively, only good industries and all industries – both goods and services. Table 2.1, for goods only, is used to determine the threshold for tradeability, while Table 2.2 illustrates how the threshold compares with the overall distribution in all industries. The comparison of Table 2.2 with Table 2.1 shows that, as expected, services industries have, on average, a higher domestic share than goods industries, reflecting that services are typically traded less than goods. Table 2.2 also shows that the threshold of 0.88 also happens to fall at exactly the 75th percentile of the distribution across all industries. Given that industries with domestic shares below 0.88 are considered tradable, this means that 75% of all industries will be tradable (recalling that 90% of goods industries were defined to be tradable).

2) Determining the ratio of the EU to global thresholds Once it has been determined which industries compete internationally, the next step is then to divide them into those competing at the EU level and those competing at the global level. The criterion used here is intuitive. Industries for which the non-EU share is larger than the EU share are defined as global; on the

Table 2.1: Industry-level summary statistics of OTT measures for goods industries only

	Mean	s.d.	min	p10	p25	p50	p75	p90	Max
OTT Domestic Share	0.63	0.19	0.12	0.41	0.54	0.65	0.78	0.88	0.98
OTT EU Share	0.20	0.11	0.01	0.07	0.12	0.19	0.27	0.38	0.46
OTT Non-EU Share	0.16	0.12	0.01	0.04	0.08	0.16	0.20	0.27	0.75
Observations	93								

Note: the table presents industry-level summary statistics on the openness to trade (OTT) measures for goods industries only.

Source: ICIO and BACI.

Table 2.2: Industry-level summary statistics of OTT measures for all industries

	Mean	s.d.	min	p10	p25	p50	p75	p90	Max
OTT Domestic Share	0.72	0.21	0.12	0.42	0.57	0.73	0.88	0.99	1.00
OTT EU Share	0.16	0.12	0.00	0.01	0.07	0.13	0.23	0.34	0.46
OTT Non-EU Share	0.12	0.12	0.00	0.01	0.04	0.09	0.18	0.25	0.75
Observations	151								

Note: the table presents industry-level summary statistics on the openness to trade (OTT) measures for all industries.

Source: ICIO, BACI, and TISP.

contrary, industries for which the EU share is larger than the non-EU share are defined as competing at the EU level.

Summary of the taxonomy

Constructing the measures of openness to trade and applying the thresholds to each industry determines the taxonomy. Each industry is assigned a unique geographical dimension on which it competes. Measures of concentration and of entrenchment are computed at the level of aggregation determined by the taxonomy, therefore providing pictures of these alternative proxies closer to competition in the relevant market. In total, there are 40 domestic industries, 85 EU industries, and 26 global industries. The full list of industries and their associated geographic market dimension are listed in Section 2.B of the Appendix (Table 2.10).

2.4.3 Concentration

Market concentration captures the share of gross output accounted for by the largest firms in a market. Its measurement involves a series of crucial methodological decisions. Specifically, it requires: defining accurately what a “relevant market” is and what “largest firms” means; measuring the boundaries of a “firm”; defining how to measure output (for both firms and industry). With respect to the existing literature, this project makes some steps ahead in the measurement of market concentration, detailed throughout this section).

The overall size of a market, $S_{s,g,t}$, is defined as the total gross output in an industry s , in its relevant geographic market g , at time t . As already mentioned, industries s are mainly defined at the 3-digit level, reaching a much higher level of disaggregation with respect to previous cross-country analyses. Geographies are defined using the taxonomy to capture the relevant geographic dimension at which competition takes place for each industry (domestic, European, global). This

work follows Bajgar, Berlingieri, et al. (2023) in defining the contribution of the leading firms to the market. Therefore, it considers the activities of business groups rather than single firms active in a given industry-geographical region (“market” in this setting). The gross output of the largest four firms is defined as $\sum_{f \in Top4} S_{f,s,g,t}$ ²³.

Therefore, the baseline Concentration Ratio (CR4) – the share of gross output accounted for by the top 4 business groups in a market - is defined as follows²⁴:

$$CR_{s,g,t}^4 = \frac{\sum_{f \in Top4} S_{f,s,g,t}}{S_{s,g,t}} \quad (2.6)$$

The remainder of this section provides detail on the methodology in computing this concentration measure.

Industry dimension

As already sketched in previous sections, when looking at concentration, economists and competition authorities try to understand what the “relevant market” is, i.e., the bundle of products and/or services that are regarded as interchangeable or substitutable by the consumer. In judging competition cases, competition authorities typically identify relatively narrow markets based on product considerations rather than on industry ones.

Although relatively straightforward from a theoretical point of view, in practice, empirical analyses are restricted by the availability of data, which is usually collected at the industry rather than at the product level. In addition, existing industry data used to construct concentration measures are normally available at a quite aggregated level, which departs from the notion of the relevant market. As a result, many existing cross-national studies rely on data at the 2-digit (or higher) industry level (for example, Bajgar, Berlingieri, et al. (2023) and Kalemli-Ozcan et al. (2023)). Studies on an individual country may have more detailed industries (such as Autor, Dorn, et al. (2020) but are unable to account for the cross-border activities of firms and the global nature of some markets. For these reasons, there have been concerns that measuring concentration at the level of industries does not accurately capture true concentration (Berry, Gaynor, and Morton (2019); Benkard, Yurukoglu, and Zhang (2021), Shapiro (2018), Werden and Froeb (2018)).

This project tries to overcome some of these concerns by developing measures of concentration for more detailed industries -typically 3-digit level, whenever possible -while still incorporating a cross country dimension and allowing for international activities of firms (see below). In total, industry concentration can be calculated for 127 industries across mining, manufacturing, utilities and non-financial market services sectors (see section Sample for Concentration and Entrenchment of Annex A for details). Manufacturing industries are almost all defined at the 3-digit level, while some services industries are slightly more aggregated to allow an accurate match with the trade

²³The numerator is obtained simply by adding the sales (as a proxy of gross output) of the largest business groups in the relative market, while the denominator is defined using the measured gross output of an industry. Using the sum of sales of all the firms contained in the Orbis dataset would not provide an accurate representation of a market in this setting since Orbis is not representative of all firms and does not capture the overall economic activity produced in a market (Bajgar, Berlingieri, et al. (2020)). Therefore, using Orbis to compute the denominator would lead to its underestimation, especially in the initial years of the sample, which in turn would introduce a downward bias in concentration trends (Bajgar, Berlingieri, et al. (2023)). Note that, for most industries, sales and gross output are very similar, so gross output will be used as a synonym for sales when looking at the numerator. However, there might be differences in certain industries, such as Wholesale and Retail. These concerns are addressed in various Robustness exercises.

²⁴Note that all over the report, “top 4 firms” refers to the four business groups (not firms) with the largest gross output in each market. The term “firm” has been preferred to “business group” for simplicity of explanation.

data (see Section 3 for details). The same set of industries is also used to define leadership ratios and entrenchment measures, described in subsequent sections.

This project tries to overcome some of these concerns by developing measures of concentration for more detailed industries -typically at a 3-digit level, whenever possible -while still incorporating a cross-country dimension and allowing for international activities of firms (see below). Industry concentration can be calculated for a total of 127 industries across mining, manufacturing, utilities and non-financial market services sectors (see section Sample for Concentration and Entrenchment of Annex A for details). Manufacturing industries are almost all defined at the 3-digit level, while some services industries are slightly more aggregated to allow an accurate match with the trade data (see Section 3 for details). The same set of industries is also used to define leadership ratios and entrenchment measures, described in subsequent sections.

Geographic dimension

The taxonomy defines the geographic dimension at which each industry competes and, therefore, over which market the proxies of competition are computed. Technological, physical, and policy-related factors, which differ for each industry, determine the geographic scope of each market. For instance, consumers in markets where it is not costly to supply products across borders can easily purchase products sourced outside their home country.

For industries that are defined to compete at the European level, the European activities of business groups are aggregated across all European countries (noting that the activities of non-European subsidiaries are excluded, as discussed in detail in the next sub-section). For example, when defining concentration, the numerator includes the gross output of the largest four firms across all European countries in each industry, and the denominator aggregates industry-level gross output across all European countries. Similarly, for industries defined as global, both firm- and industry-level gross output are aggregated over all countries in the sample²⁵. In contrast, the concentration ratio is computed within each country-industry for industries defined to compete domestically.

An advantage of aggregating tradable industries across countries is that it accounts for all trade between countries within the region. For example, in global markets, the entire activities of firms in each industry are accounted for, regardless of where production and consumption occur (although note that with the data limitations, some countries remain excluded). By aggregating firm- and industry-level gross output across countries within their geography of competition, it is not necessary to make any adjustments for trade within the region. This result is a very relevant point since firm-level data on international trade are not available in the Orbis database (and are almost nonexistent in a cross-country setting). It implies that it is not possible to identify exactly, for each firm, the share of gross output to be assigned in the various countries where the firm is active; as a consequence, they are all assigned to the country where it is located. Therefore, the taxonomy provides a conceptually important and empirically practical solution to account for the globalised nature of highly tradable industries in computing concentration²⁶. A later sub-section discusses the incorporation of international trade in the measures of concentration, in addition to the use of the taxonomy.

²⁵Recall that, outside the EU, data are available for three additional countries: Korea, Japan, and USA.

²⁶In principle, when looking for example at industries competing at the domestic level, only the gross output of the top firms sold domestically should be considered in the numerator of concentration, subtracting, therefore, their exports. On the contrary, when looking at industries competing at the European level, all exports of the top 4 firms to other European countries are correctly included in the numerator of the CR4.

Identifying activities of business groups

Large firms operating in the same industry may not be independent of each other but rather be part of the same business group (Altomonte et al. (2021)). This is likely to be even more relevant when concentration is measured at the global or European level since business groups often serve each country through a different firm entity. For example, in the sample of this report, the average “top 4” business group has seven subsidiaries within the same industry-region. Therefore, it is crucial to account for the gross output of business groups when measuring concentration and not to focus on individual firms or any other economic entity.

Following the work of Bajgar, Berlingieri, et al. (2023) and Bajgar, Berlingieri, et al. (2019) (who, in turn, build on Bloom, Schankerman, and Reenen (2013)), this report focuses on the activities of business groups. When measuring concentration, different subsidiaries belonging to the same business group and active in the same market are treated as if they were a unique entity, as neglecting these ownership linkages may lead to an understatement of concentration. Bajgar, Berlingieri, et al. (2023) show indeed that considering this business group dimension has a substantial impact on the resulting concentration measures. Therefore, accounting for linkages between firms in the same business group is of utmost importance when measuring concentration.

To construct the relevant activities of a business group in each market, the unconsolidated gross output of each subsidiary is aggregated across all subsidiaries that operate in the relevant geography and industry^{27,28}.

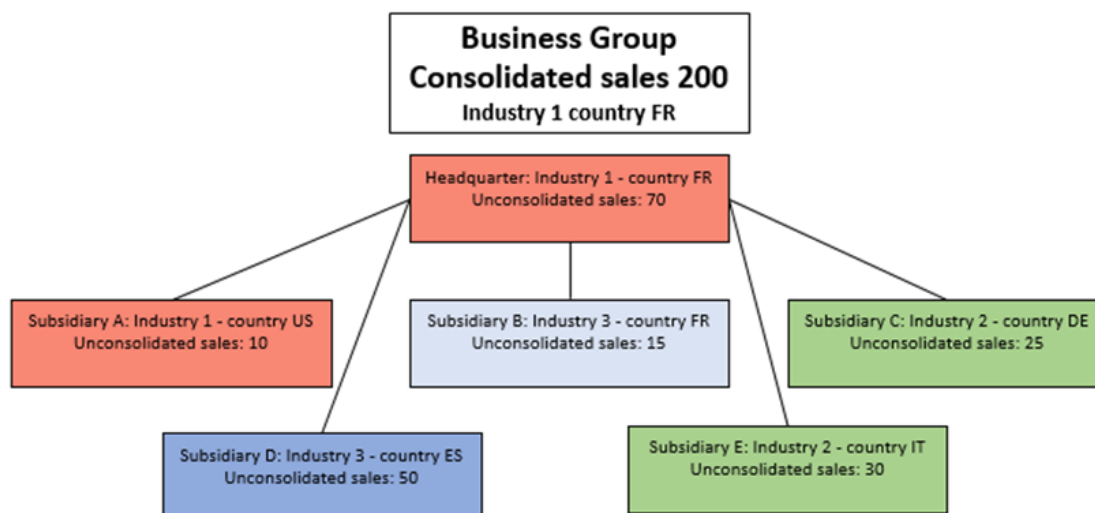
To provide an intuition of the methodology proposed aimed at linking business groups’ activities among different markets, Figure 2.3 provides an illustrative example of a business group composed of i) the headquarter and ii) five subsidiaries, operating in three different sectors and five different countries. The first information exploited is the unconsolidated financial information of subsidiaries, obtained from Orbis, and aggregated according to the ownership information from the Orbis-Zephyr database. Then, for each industry, the unconsolidated gross output is aggregated at the geographic level specified by the taxonomy. This apportioning procedure allows to split the activities of the same business group across industries-geographies pairs (markets). Returning to the example, assume that the market for Industry 1 is global, for Industry 2 is European, and for Industry 3 is domestic. Then the business group is active in four different markets (each represented by a different colour in the figure): i) in industry 1, which competes at the global level, with a total gross output equal to 80 (headquarter 70 + subsidiary A 10); ii) in industry 2, which competes at the European level, with a total gross output equal to 55 (subsidiary C 25 + subsidiary E 30); iii) in industry 3 in Spain, which competes at domestic level, with total gross

²⁷Consolidated accounts are used only to correct the unconsolidated information in cases where the total subsidiary sales exceed group sales (due to inter-company transactions) or where unconsolidated data are missing. See Annex A for further details. Additionally, consolidated accounts are used directly only in two cases in which consolidated and unconsolidated accounts coincide: for independent firms (firms that are not part of a business group) and for subsidiaries at the bottom of the ownership hierarchy (subsidiaries not owning any other subsidiary) that do not report unconsolidated accounts.

²⁸With the exception of Bajgar, Berlingieri, et al. (2023), previous studies have mainly followed three approaches to deal with cross-ownership linkages. A first approach is to neglect business groups and focus only on unconsolidated information of individual firms. This method underestimates the concentration if multiple firms in the same market are part of the same group. A second approach is to neglect subsidiaries and focus only on the consolidated accounts of the headquarters. This method attributes the entire activity of the business group to the headquarters, overestimating (underestimating) the concentration in the headquarters’ (subsidiaries’) market. A third approach is to include both the activity of the business group and the firm subsidiaries and try to address the issue of double counting by dropping the unconsolidated information for headquarters -the most obvious source of double counting -but still including the unconsolidated information of subsidiaries. This method overestimates concentration since it double counts subsidiaries’ revenues. See Bajgar, Berlingieri, et al. (2019) for additional explanations on these alternative approaches.

output equal to 55 (subsidiary D); and iv) in industry 3 in France, with total gross output equal to 15 (subsidiary B). These are the total gross output that, in the example, would be used to compare the total gross output of this group with that of other groups/firms competing in the same market. For further details, see Annex A and Bajgar, Berlingieri, et al. (2023), (2019[16]).

Figure 2.3: Example of the apportioning technique of Business Group activities



Note: figure depicts an example of a hypothetical group consisting of a parent company from France and operating: i) in industry 1 (global), with a US subsidiary in the same industry; ii) in industry 2 (European) with two subsidiaries, one from Germany and one from Italy; and finally iii) in industry 3 (domestic) with a subsidiary in France and one in Spain. The different colours identify the 4 different markets in which the group is active: the global market in industry 1 (with total gross output of 80), the European market in industry 2 (with total gross output of 55), the French market in industry 3 (with total gross output of 15), and finally the Spanish market in industry 3 (with total gross output of 50).

Effective market size: accounting for international trade

While the taxonomy accounts for international trade between countries within a region, it does not account for trade from outside the region (and from countries not included in the sample). Trade can affect both the numerator and denominator of concentration. On the denominator, imports increase the overall size of a market, while the value of goods exported by producers could be deducted from the market size. Similarly, the gross output of top firms going outside of the region could be subtracted from the numerator, or as an alternative, foreign firms serving a market only through imports could be among the top firms in the numerator. Amiti and Heise (2021) show that, in the US, industry concentration is flat when accounting for the gross output of foreign exporters, while it is increasing if the import correction is not made.

However, in this setting, it is important to note that there is an interaction between the taxonomy and the import correction. The taxonomy is defined using data on international trade. Industries with high values of trade will be either European or global. Therefore, only the imports from outside the region are added to the denominator. For domestic industries, imports comprise a smaller share of production. Hence, in the presence of the taxonomy, the import correction is less

impactful²⁹. Despite this, as a robustness check, international trade from outside a region g is incorporated into the concentration measures. The adjusted measure can be written as:

$$CR_{s,g,t}^4 = \frac{\sum_{f \in Top4} S_{f,s,g,t} - \alpha X_{s,g,t}^{c \neq r}}{S_{s,g,t} + M_{s,g,t}^{c \neq r} - X_{s,g,t}^{c \neq r}} \quad (2.7)$$

The denominator is adjusted to account for imports into the region ($M_{s,g,t}^{c \neq r}$) and exports from the region ($X_{s,g,t}^{c \neq r}$). Note that imports and exports between countries within a region, such as the gross output of Italian firms in Spain for an industry defined to compete at the European level, do not need to be corrected as they are accounted for in the production value of the exporter (Italy, in this example)³⁰.

A further correction can be made in the numerator to account for the value of the top 4 firms' gross output that is exported to different markets. However, as firm-level data on trade are not available, an assumption must be made on the share of total exports that is accounted for by the top 4 firms, denoted by α . For example, α could be equal to the share of total production accounted for by the top 4 firms, although this is likely to underestimate their share as larger firms are more likely to export (Bernard et al. (2012))³¹. These robustness checks are implemented in two ways. The first only makes a correction for imports in the denominator, while the second makes a correction for both imports and exports (in both the denominator and numerator).

Leadership ratio

CR4 provides information on the share of the top 4 firms, but does not reveal whether, for instance, the market is monopolistic or oligopolistic. In monopolist markets, the market share of the leading firm (top 1 in each market) significantly outweighs that of the followers, while in oligopolistic markets, the market shares of the top two firms are high but similar. These market structures can have contrasting implications for consumer prices and choice sets, as well as market dynamics and policy design.

Two measures of leadership ratios are computed to explore the relative market shares of the leading firm relative to the following firms. These measures are defined using the information on the top 4 firms included in the numerator of CR4.

The first measure is the 2-firm leadership ratio, $LR_{s,g,t}^{2-firm}$, defined as the gross output of the leading firm over the gross output of the second leading firm within a market:

$$LR_{s,g,t}^{2-firm} = \frac{GO_{s,g,t}^{First}}{GO_{s,g,t}^{Second}}. \quad (2.8)$$

When the leadership ratio is, for example, equal to 2, the leading firm has a gross output twice as big as the gross output of the second biggest. A high leadership ratio is associated with a highly

²⁹An important caveat is that, due to data availability, there are only three additional countries (Japan, Korea, and USA) included in the global category. Therefore, the import correction always accounts for imports from other non-EU countries, such as China.

³⁰Note that the sample of countries is not complete, so imports from countries not included in the sample are also incorporated in the denominator.

³¹Note that, unlike Amiti and Heise (2021), firm-level international trade data is not available in the present study. If the foreign exporters ranked in the top 4 firms, then the measure of concentration would be downward biased. Amiti and Heise (2021) show that exporters into the US market tend to be smaller, which explains why accounting for increases in imports flattens the concentration trend.

monopolistic market, although note that the leadership ratio can be large even if the market share of the biggest firms is relatively small (i.e., if the leader, although not representing a big share of the market, is much bigger than the second one). For the leadership ratio measure to be meaningful, it is therefore important to couple it with the concentration one.

The second measure is the 4-firm leadership ratio, $LR_{s,g,t}^{A-firm}$, measured as the gross output of the first firm over the sum of the gross output of the second, third and fourth firm in the ranking:

$$LR_{s,g,t}^{A-firm} = \frac{GO_{s,g,t}^{First}}{GO_{s,g,t}^{Second} + GO_{s,g,t}^{Third} + GO_{s,g,t}^{Fourth}}. \quad (2.9)$$

If the ratio is larger than one, it means that the first firm has a gross output bigger than the sum of the gross output of the other three. In contrast, if the ratio is smaller than one, it means that the gross output of the leading firm is smaller than the sum of the gross output of the other three. The market has more competition among the top 4 when the ratio is below one. This situation is compatible with two settings: if the shares of each of the four leading firms are large, then the market is oligopolistic, whereas if each share is small, then the market may be competitive (at least in terms of market shares).

2.4.4 Entrenchment

Both the concentration ratio and the leadership ratio are static measures, as they consider the market shares of the leaders at each point in time but do not follow them over time. However, market dynamism is an important feature of competitive markets. Very concentrated markets could still be contestable if firms at the top are competing to get the leadership such that the identity of the market leader changes over time. On the contrary, in industries that are less competitive, top firms may be more entrenched. That is, leading firms can remain persistently as market leaders over the long term, with negative consequences for competition.

To measure the extent of churning among top firms in each industry, the persistence of firms in the group of the “top 4” is calculated. This entrenchment measure aims at capturing the likelihood that market-leading firms remain as market leaders between two consecutive periods³².

Specifically, market entrenchment is computed as the number of firms in the top 4 at time t that were also in the top 4 at time $t-1$ in each market. In each year, if all the firms in the top 4 were in the top 4 in the previous year, then the entrenchment measure is equal to 4. On the contrary, if none of the firms in the top 4 in t were in the top 4 in $t-1$, the measure equals zero. Hence, the entrenchment measure is bounded between 0 and 4. Note that it is closely related to that of Competition and Authority (2022) and Bajgar, Criscuolo, and Timmis (2021), and it can be written as:

$$Ent_{s,g,t} = \sum \{f \in Top4\} \wedge \{f \in Top4_{t-1}\} \mathbb{1}_{f,s,g,t} \quad (2.10)$$

where $\mathbb{1}_{f,s,g,t}$ is an indicator equal to one if business group f , active in industry s and geography g at time t , was also among the top 4 in the previous year and in the same market. To check the robustness of the results, the same measure has also been computed for different time horizons:

³²In line with the concentration measures, to define the top 4 firms with the largest gross output in each market the business group level information has been considered.

between t and $t-2$, and between t and $t-3$. In this case, a business group is defined as entrenched in the top 4 if it is in the top 4 across all years of the time interval considered³³.

³³For example, in the entrenchment measure over three consecutive years, the measure captures the number of firms that are the top 4 in t , $t-1$, and $t-2$. As a further robustness check, a measure of entrenchment based on the methodology developed by Bessen et al. (2020), which considers the displacement hazard, is also constructed.

2.5 Trends

2.5.1 Concentration

As already mentioned, concentration is measured as the share of gross output accounted for by the four largest firms in a market. A market is defined using the taxonomy, which assigns each industry to a unique “geographical bucket”: domestic, European, or global.

In the report, concentration trends are presented separately for each geographical bucket. The evolution of the baseline concentration measure is presented by plotting both levels and the cumulative unweighted average change since the year 2000 (normalised to 0) in each geographical bucket. When looking at the cumulative change, in each year, the average yearly change is computed across all industries within a bucket for the European and the global buckets and across all country-industries pairs within the domestic one. Then, the overall cumulated change is computed and plotted by summing up the yearly changes starting from the base year 2000³⁴.


The unweighted average is chosen as a baseline for three reasons. First, when looking at overall concentration trends in each bucket, it is preferable to give all industries the same weight to avoid a representation informative only of a few big industries. Second and related, due to data limitations, there is substantial heterogeneity in the level of aggregation of industries, and more aggregated industries would mechanically get more weight in the weighted trends. As a result, trends would potentially be driven by few big industries and, as such, not very informative of the overall evolution of concentration. Finally, once the economy is divided into geographic buckets, the weight assigned to each industry within the bucket would not coincide with the share of that industry in the overall economy and, thus, with its relative importance. However, weighted trends are discussed later and produced as a robustness check in Annex C. In addition, this section provides a qualitative discussion of the heterogeneous evolution of concentration across different countries and industries.

Figure 2.4 and Figure 2.5 report concentration levels and cumulative unweighted average changes across the three geographical buckets. The average concentration level is higher in industries competing at the domestic level than in those competing at the European and global levels, which show similar levels of concentration. Specifically, the top 4 firms represent, on average (over the period considered), around 43% of the total gross output of the industry in industries competing at the domestic level, and around 26% in the other geographical buckets. In addition, Figure 2.5 shows that, on average, industry concentration has increased across all geographic buckets. Industries that compete at the domestic level had the greatest increase in average concentration, by around 6 percentage points (p.p., henceforth) between 2000 and 2019. Industries that compete internationally – either at the European or global level – increased their concentration by approximately 4 p.p. However, while industries competing at the domestic and the European level see a smooth increase over the period considered, in industries competing at the global level concentration is relatively stable up to about 2012 and then starts to slightly increase³⁵.

³⁴Concentration trends are presented both in levels and by means of cumulative growth. Note, however, that given that the sample used in the analysis is not fully balanced, the cumulative average changes allow to control for any change in sample composition over time, while trends in levels can potentially exhibit jumps if any of the industries entering/exiting has substantially different levels of concentration with respect to the average level. Concentration levels should always be considered with caution due to data limitations. Levels of concentration might be sensible to specific data issues for various reasons, such as: different data sources and definitions of the main variables for the numerator and the denominator; missing ownership links, with a resulting higher value for the numerator; intragroup sales not always captured; output volatile and difficult to measure at the 3-digit in certain industries. Therefore, trends of cumulative changes are more stable and, as such, are the preferred option in the report to show the evolution of concentration over the years.

³⁵Note that the global bucket includes only 20 industries, while the European one comprises 80 industries and the

Figure 2.4: Concentration levels across geographical buckets




Note: chart shows the unweighted average across industries (and countries, for the domestic bucket) of CR4 levels. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are classified as either domestic, European, or global, depending on the results of the taxonomy. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE for the domestic and European bucket, while in the global one also JAP, KOR, and USA are present.

Source: OECD calculations.

domestic one 27 industries (*15 countries). Therefore, the global trend is slightly more volatile than the other buckets, as it includes fewer observations and, as such, it is more sensible to single industries' changes in concentration.

Figure 2.5: Concentration cumulative changes across geographical buckets



Note: the chart shows the unweighted average across industries (and countries, for the domestic bucket) of CR4 cumulative growth. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are classified as either domestic, European, or global, depending on the results of the taxonomy. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE for the domestic and European bucket, while in the global one also JAP, KOR, and USA are present.

Source: OECD calculations.

The aggregate trends reported in Figure 2.4 and Figure 2.5 can hide substantial heterogeneity across industries (and countries, for the domestic bucket), and therefore extensive sensitivity tests have been performed to understand whether these trends are driven by few observations. The main exercise performed to study heterogeneity across sectors consisted in removing industries from the sample one at the time to check their relative importance in driving the cumulative changes trends³⁶. In addition, for industries competing at the domestic level, where also the country dimension can be investigated, a similar exercise has been conducted by removing one country at the time³⁷. All in all, the trends appear to be robust to the exclusion of single industries and countries, as no single countries or industries drive the cumulative changes of concentration in any

³⁶Domestic bucket: only three industries marginally affect the average cumulative change trend when removed, either by decreasing it by 1 p.p. (091, Support activities for petroleum and natural gas extraction; 352, Manufacture of gas, distribution of gaseous fuels through mains) or by increasing it, again, by 1 p.p. (353, Steam and air conditioning supply). European bucket: there are four industries that make the cumulative change in concentration increase by about 1 p.p. each when removed from the sample (232, Manufacture of refractory products; 242, Manufacture of tubes, pipes, hollow profiles and related fittings, of steel; 781 Activities of employment placement agencies; 783 Other human resources provision) and one that make the cumulative change in concentration decrease by almost 2.5 p.p. (262, Manufacture of computers and peripheral equipment). While the variation in the last industry appears to be relevant, in terms of the average cumulative change in concentration the net effect is essentially nihil when considering also the former four. Global bucket: there are only two industries marginally driving the trends, in two opposite directions. Dropping from the sample industry 151 (Tanning and dressing of leather; manufacture of luggage, handbags, saddlery, and harness; dressing and dyeing of fur) would make the cumulative change of the global bucket decrease by about 1.5 p.p., while dropping industry 303 (Manufacture of air and spacecraft and related machinery) would make the trend increase by about 2 p.p. Note that single data points should be interpreted with caution, as explained in footnote 34.

³⁷Only Poland, Portugal, Sweden, and the UK can be considered as marginally driving the overall growth in the domestic bucket (Poland and Sweden both decrease the cumulative growth by about 1 p.p. each, while for Portugal and the UK the opposite is true; overall the effect of these countries on the aggregate growth cancels out).

of the three geographical buckets.

Further discussion and robustness checks

Several exercises are performed and briefly described in this subsection to assess the robustness of the results and further explore aspects of market concentration. In particular, additional evidence is reported on the role of weighting in computing aggregate concentration measures, on trade adjustments, and on the evolution of the distribution of concentration.

Figure 2.20 and Figure 2.21 in Section 2.B.1 of the Appendix shows, respectively, concentration levels and cumulative changes weighted by market gross output (country-industry for the domestic bucket, industry for the European and global buckets) within each geographical bucket. The comparison of unweighted levels of concentration (Figure 2.4) with the weighted ones (Figure 2.20) reveals that weighting by market size (in terms of gross output) reduces the level of aggregate domestic concentration, suggesting that concentration is higher in smaller markets (country-industries pairs in this case). On the contrary, the weighting procedure increases the aggregate level in the global bucket and, to a lower extent, also in the European bucket, indicating that concentration is higher in relatively bigger industries. In addition, the comparison of unweighted and weighted concentration cumulative changes (respectively, Figure 2.5 and Figure 2.21) shows that when weighting for the relative importance of the markets, concentration looks overall flat over the period 2000-2019 for industries competing both domestically and at the European level. In industries competing at the global level, the weighted trend is even decreasing. The combination of these results suggests that for the domestic and European buckets, the increase in concentration mostly occurs in relatively small markets (in terms of gross output), while for the global buckets, the decrease occurs in relatively big sectors.

As explained in Section 2.4, concentration measure crucially depends on the definition of the market. Amiti and Heise (2021) show that accounting for the trade interconnections, which allows us to consider the actual size of a market, significantly affects concentration trends. Following this seminal paper, a robustness check that accounts for the role of import and export is performed (see Section 4 for details). In the next three figures (Figure 2.6 for domestic markets, Figure 2.7 for European markets and Figure 2.8 for global markets respectively): i) the blue line reports the baseline with no correction for international trade (same as in Figure 2.4); ii) the orange line adds the import correction in the denominator, but no export correction; iii) the green line not only accounts for imports but also exports, which are subtracted from both the denominator and the numerator. Specifically, in the import correction exercise, imports in the same industry from outside the geographical bucket considered (domestic, European, global) are subtracted. In the exercises in which exports are also accounted for, exports in the same industry to the rest of the world (with respect to the bucket considered) are fully subtracted from the denominator (to account only for the gross output that is consumed in each market), while at the numerator it is subtracted a share of exports corresponding to the share of gross output accounted for by the four largest firms in the market (since data do not provide information on the export flows at the firm level).

In line with Amiti and Heise (2021), the trade corrections tend to dampen the rise in concentration. As largely expected, trade corrections affect mostly the trends of cumulative changes in the European and global buckets, which are more tradeable, whereas the growth of concentration in the domestic bucket remains virtually unchanged.

Figure 2.6: Trade adjustments, industries competing domestically



Note: the chart shows the unweighted average across industries and countries of cumulative change in CR4 in the domestic geographical bucket for different types of trade adjustments. The blue line refers to the baseline average cumulative change (no corrections). The orange one to the correction obtained by adding import at the denominator. The green one to the correction obtained by adding import and subtracting total exports at the denominator, and by subtracting at the numerator a fraction of export given by the share of gross output accounted for by the four largest firms in the market. Industries are a mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.


Figure 2.7: Trade adjustments, industries competing at the European level



Note: the chart shows the unweighted average across industries and countries of cumulative change in CR4 in the European geographical bucket for different types of trade adjustments. The blue line refers to the baseline average cumulative change (no corrections). The orange one to the correction obtained by adding import at the denominator. The green one to the correction obtained by adding import and subtracting total exports at the denominator, and by subtracting at the numerator a fraction of export given by the share of gross output accounted for by the four largest firms in the market. Industries are a mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.

Figure 2.8: Trade adjustments, industries competing at the global level



Note: the chart shows the unweighted average across industries and countries of cumulative change in CR4 in the global geographical bucket for different types of trade adjustments. The blue line refers to the baseline average cumulative change (no corrections). The orange one to the correction obtained by adding import at the denominator. The green one to the correction obtained by adding import and subtracting total exports at the denominator, and by subtracting at the numerator a fraction of export given by the share of gross output accounted for by the four largest firms in the market. Industries are a mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, JAP, KOR, NOR, POL, PRT, SVN, SWE, and USA.

Source: OECD calculations.

As an additional robustness, Figure 2.22 and Figure 2.23 report the average concentration level and the average cumulative change of concentration when all industries are treated as competing at the European level (i.e., concentration is computed as if all industries were belonging to the European geographical bucket): in this case, there is an overall increase in concentration of about 5 p.p.. This result is in line with Bajgar, Berlingieri, et al. (2023), which found an average increase of concentration at the European level of about 3 p.p. over the years 2000-2014³⁸.


While aggregate trends are informative of what happens on average in the economy, it is also worth investigating how the distribution of concentration evolved throughout the years of the sample (2000-2019). From an economic perspective, it is indeed important to understand whether the overall increase in concentration comes from industries that had already high levels of concentration, or low levels, or rather it is driven by industries with different levels of concentration (see, for example, Koltay, Lorincz, and Valletti (2022)). To perform this exercise, concentration has been analysed in different deciles of its distribution. In order to meaningfully do so, it is important to have a high number of industries. Thus, the exercise is carried out looking at concentration measures assuming that all industries belong to the European geographical bucket³⁹. Specifically,

³⁸Please notice that the level of aggregation, the sample of countries, and the concentration measure (largest eight business groups instead of largest four) adopted in Bajgar, Berlingieri, et al. (2023) are different from those used in this work.

³⁹Maintaining the taxonomy would imply having buckets with too few industries to meaningfully compute concentration deciles within each of them (notably in the global bucket). Moreover, simply pooling industries while maintaining the taxonomy would possibly distort the results since, as discussed above, industries belonging to the domestic bucket have an average concentration level that is higher than those belonging to the global and the

industries are grouped year by year in quantiles based on their concentration level⁴⁰. Figure 2.9 shows the average level of selected quantiles of the concentration distribution and Figure 2.10 reports the cumulative changes of the average level of such quantiles⁴¹. First, Figure 2.9 shows that the top decile exhibits an average level of concentration (over the period considered) that is more than two times that of the 50th-90th quantile (85% vs. 34%) and more than seven times that of the bottom half of the concentration distribution (85% vs. 11%). In addition, Figure 2.10 shows that the top decile is the one that has been growing the most (about 25 p.p.) along the period considered, while other quantiles grow at a slower pace (2.5 p.p. on average).

Figure 2.9: Concentration levels across selected quantiles



Note: the chart shows the average concentration level -smoothed through a moving average filter- of selected quantiles of the concentration distribution. All industries are considered as belonging to the European bucket. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.

European bucket (see Figure 2.4 and related discussion).

⁴⁰Creating deciles year by year implies that the same industry can be in different deciles over the years (e.g., industry 182, “Recorded media”, belongs to the 10th decile in 2018 and in the ninth decile in 2017).

⁴¹The average level is smoothed through a moving average filter with two leads and two lags. This choice is due to a certain degree of variability in the average concentration level of the top decile in the initial years. The filter applied does not substantially modify the average level of concentration in the initial and final years. As a consequence, the smoothed cumulative change (which is computed by taking the cumulative growth of the smoothed average level) exhibits a more regular pattern while preserving the same magnitude of the non-smoothed trend.

Figure 2.10: Cumulative changes across selected quantiles



Note: the chart shows the cumulative changes in concentration level of selected quantiles of the concentration distribution. All industries are considered as belonging to the European bucket. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.

Notably, while these charts provide evidence for the evolution of the average level of single quantiles, they remain silent about their composition in terms of industries. Figure 2.24 provides information on the number of industries that entered at least once in each decile and suggests that there is a considerable amount of churning, especially in the central deciles of the distribution. For example, 22 distinct industries (out of 127) appear in the first decile throughout the time span of the sample, 59 enter at least once in the fifth one. Focusing on the top decile, Figure A C.6 lists the 32 industries that appear at least once in the top decile. It emerges that, on average, every year, 2.2 new industries enter it. Each industry stays in the top decile for 7.5 years on average, even if also in this case there is substantial heterogeneity, with some industries always present and others entering only for one year. The overall interpretation of this exercise is that the increase in concentration showed in Figure 2.25 (slightly less than 5 p.p.) seems to be driven by the increase in the average level of concentration of industries belonging to the top decile, similarly to Koltay, Lorincz, and Valletti (2022).

Leadership ratio

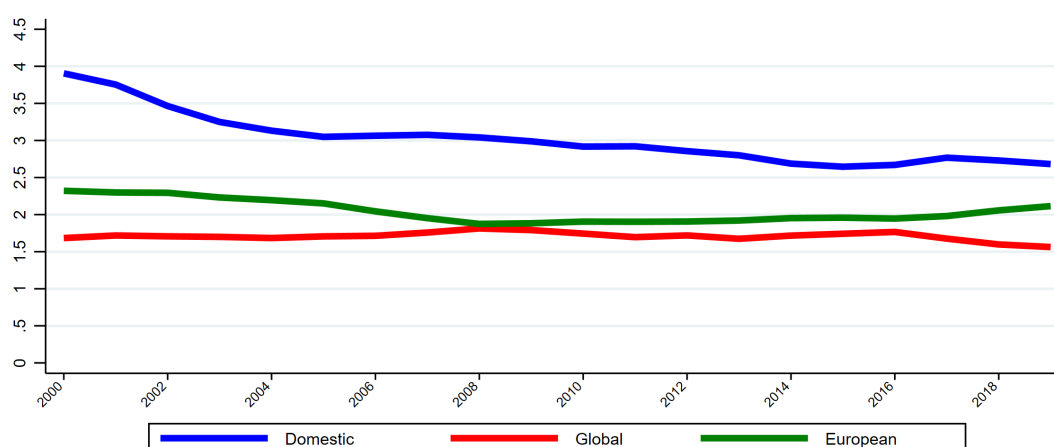
The concentration ratio measures the average market share of the top 4 firms, but it does not reveal information about the relative size of the leading firm with respect to the followers. The leadership ratio is defined as the gross output of the leading firm over the gross output of the second leading firm within a market or, in an alternative version of the measure, as the gross output of the leading firm over the gross output of the three next largest firms.

Figure 2.11 plots the moving average of the unweighted leadership ratio by geographical bucket⁴².

⁴²The reason for using the moving average is that the measure appears to be more volatile than the other concentration measures. This is because it is based on only two observations at the business group level for each market and, therefore, is more prone to firms' sales volatility and fluctuations in the business cycle. For each year t , the moving average computation uses three years ($t - 1$, t and $t + 1$, that is, one year lag, the actual year, and

There are interesting differences in the average level of the leadership ratio in tradable (European and global) versus non-tradable industries. Industries that compete at the national level have a significantly higher average leadership ratio than industries competing internationally. In 2000, in industries competing domestically, the leading firm is, on average, 3.9 times larger than the following one, whereas in industries competing at the European and global levels, the leading firm is, respectively, 2.4 and 1.6 times bigger than the following one. Interestingly, this difference blunts over the period considered: in industries competing domestically, the average leadership ratio falls over time, from 3.9 in 2000 to 2.7 in 2019; in contrast, in both European and global industries, the trend remains relatively stable over the years. Therefore, in industries competing domestically, there seems to be increasing competition between the top 2 firms over the years, but the differences in terms of gross output among them remain still larger compared to industries competing at the European and global levels.

Figure 2.11: Leadership ratio across geographical buckets (1st over 2nd), unweighted



Note: the chart shows the unweighted average across industries (and countries, for the domestic bucket) of leadership ratio considering the ratio of the sales of the leading firm over the sales of the second firm, between 2000 and 2019. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are classified as either domestic, European, or global, depending on the results of the taxonomy. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE for the domestic and European bucket, while in the global one also JAP, KOR, and USA are present.

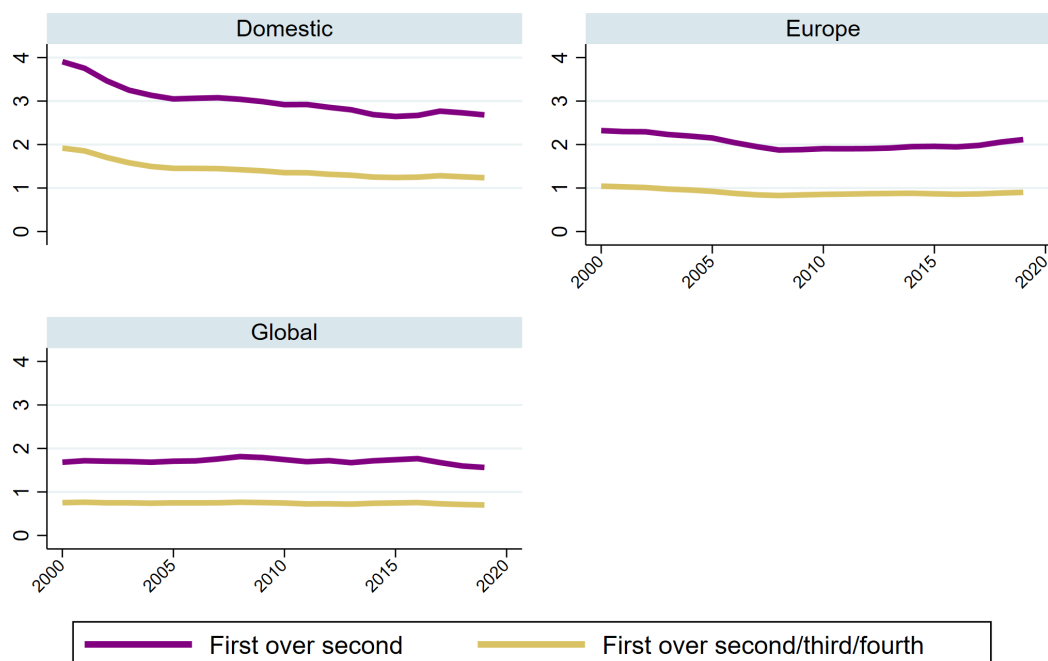
Source: OECD calculations.

Figure 2.12 shows the leadership ratio presented before alongside the alternative specification of the measure, which is computed as the gross output of the first firm over the sum of the sales of the second, third and fourth firms within each industry. In general, the alternative measure follows a very similar pattern to the previous one. In industries competing domestically, the largest firm is about 1.5 times larger than the other three firms together, with the gap having decreased over time. In industries competing at the European and global level, instead, the largest firm has about the same level (or slightly less for industries competing globally) of gross output as the sum of the other firms in the top 4.

one year lead) and takes the average among them. For the first (last) year of the time series, only the lead (lag) is considered.

These trends are broadly similar when computing the weighted average instead of the unweighted one, as shown in Figure 2.26. In addition, the unweighted trends are also qualitatively robust to the exclusion of individual industries⁴³.

Figure 2.12: Leadership ratio: alternative measures



Note: the chart shows the unweighted average across industries (and countries, for the domestic bucket) of leadership ratio considering the ratio of the sales of the leading firm over the sales of the second firm versus the ratio of the sales of the first firm over the sum of the sales of the second, third and fourths firm in the ranking, between 2000 and 2019. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are classified as either domestic, European, or global, depending on the results of the taxonomy. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE for the domestic and European bucket, while in the global one also JAP, KOR, and USA are present.

Source: OECD calculations.

2.5.2 Entrenchment

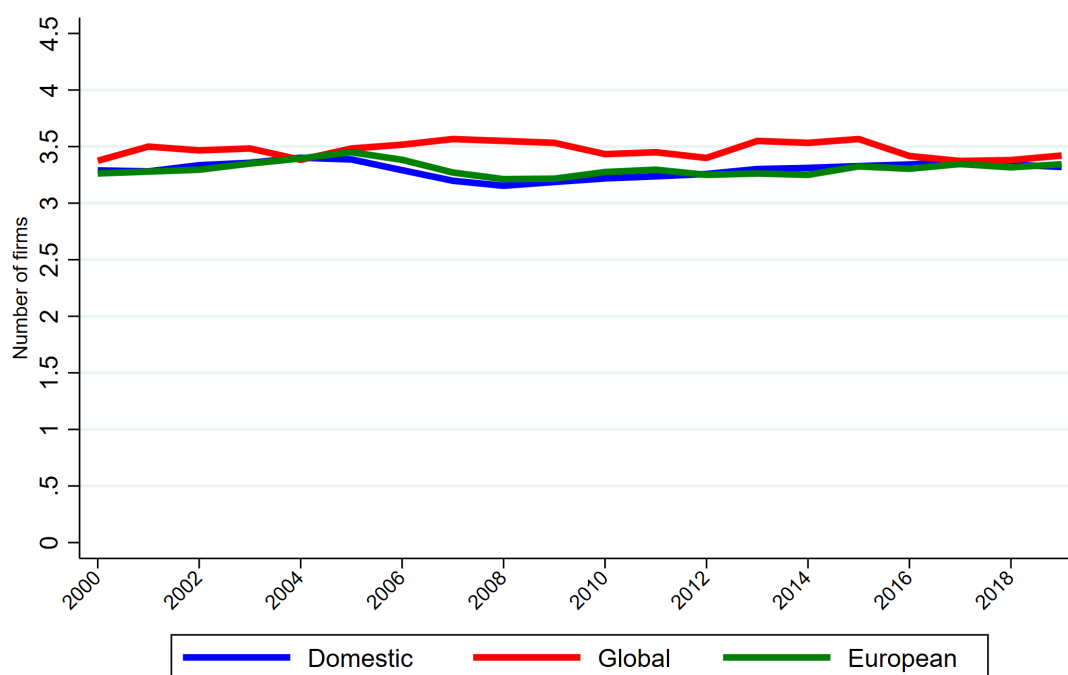
The entrenchment of top firms as market leaders is an important indicator of market dynamism. Industries may be very concentrated but simultaneously have competition among the market leaders. As shown in the methodology section, the baseline entrenchment measure is defined as the

⁴³While individual industries do not drive the evolution of the trends, the levels of the leadership ratio can, of course, change with the exclusion of some specific industries. For example, the leadership ratio trend in the global bucket is partially driven by two industries: “Manufacture of musical instruments” and “Manufacture of weapons and ammunition”, which raise the leadership ratio. In the European bucket, instead, “Manufacture of other general-purpose machinery” increases the leadership ratio even if it does not affect the trend. In the domestic bucket, the leadership ratio trend is not driven by individual markets.

number of firms that were in the top 4 in year $t - 1$ and remain in the top 4 in t in each market. As such, it is bounded between 0 and 4.

Figure 2.13 plots the yearly unweighted average of entrenchment across industries in each geographical bucket⁴⁴. On average, more than three firms that were in the top 4 in $t - 1$ remain in the top 4 in t . Overall, there is a high level of persistence in the entrenchment rate in all the geographical buckets. If anything, industries competing at the global level have, on average, slightly a higher persistence compared to industries competing at the domestic and European levels. In addition, for industries competing at the domestic and the European level, the entrenchment measure slightly decreased between 2005 and 2008, to start increasing again after 2008 to the pre-2005 levels.

Figure 2.13: Entrenchment across geographical buckets



Note: the chart shows the unweighted average across industries (and countries, for the domestic bucket) of entrenchment in the top 4 firms between 2000 and 2019. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are classified as either domestic, European, or global, depending on the results of the taxonomy. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE for the domestic and European bucket, while in the global one also JAP, KOR, and USA are present.


Source: OECD calculations.

As a robustness, Figure 2.14 compares entrenchment measures computed for different time horizons: one year before ($t - 1$) (i.e., the baseline plotted in Figure 2.13), two years before ($t - 2$), and three years before ($t - 3$). In this exercise, a firm is considered entrenched in the top 4 if it is in the top 4 across all years of the time interval considered.

⁴⁴The moving average is computed in the same way as in leadership ratio.

The three measures follow very similar patterns in each geographical bucket. Entrenchment is high and relatively flat over the period for all the different time horizons. On average, more than two firms are consistently in the top 4 every year over a three-year time horizon.

Figure 2.14: Entrenchment for different time horizons



Note: the chart shows the unweighted average across industries (and countries, for the domestic bucket) of entrenchment for different time horizons (one year before ($t-1$), two years before ($t-2$), and three years before ($t-3$)), between 2000 and 2019. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are classified as either domestic, European, or global, depending on the results of the taxonomy. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE for the domestic and European bucket, while in the global one also JAP, KOR, and USA are present.


Source: OECD calculations.

These trends are not driven by any single industry and exhibit only minor changes when industries are removed one at a time. As a robustness, weighted entrenchment trends are computed but not reported: trends are very similar to the unweighted ones. Furthermore, for the measures of entrenchment over two or three years, an alternative specification is considered whereby the firm is only required to be in the sample in the initial period ($t-2$ or $t-3$), respectively, and in the final one (t). Results (not reported for brevity but available upon request) are consistent with this alternative specification. Finally, the results are consistent when measuring entrenchment with a measure that follows Bessen et al. (2020)⁴⁵.

⁴⁵Conceptually, entrenchment and the measure used in Bessen et al. (2020) are very similar. The latter, being a hazard function, also considers the probability that firms 'entry and exit and, as such, to be reliable must be built on a sample with the population of firms, while entrenchment measures only require information on the largest four

A complementary perspective on entrenchment can be drawn by dividing markets into five categories, according to the number of new firms entering the top 4 group every year⁴⁶. Then, the share of markets falling in any of these five categories is computed for each year within each geographical bucket. Figure 2.15 shows that more than on average 80% of the markets have either no entrance of new firms in the top 4 from one year to the other (about 40%) or just one new firm (an additional 40%) in the top 4, in all geographical buckets. Importantly, global markets have the lowest share of new firms entering in the top 4.

Figure 2.15: Share of markets by new top4 firms within each geographical bucket



Note: the chart shows the share of markets between 2000 and 2019 within industries (and countries, for the domestic bucket) by five categories: no new firms in the top 4, one new firm in the top 4, two new firms in the top 4, three new firms in the top 4, and four new firms in the top 4. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are classified as either domestic, European, or global, depending on the results of the taxonomy. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE for the domestic and European bucket, while in the global one also JAP, KOR, and USA are present.

Source: OECD calculations.

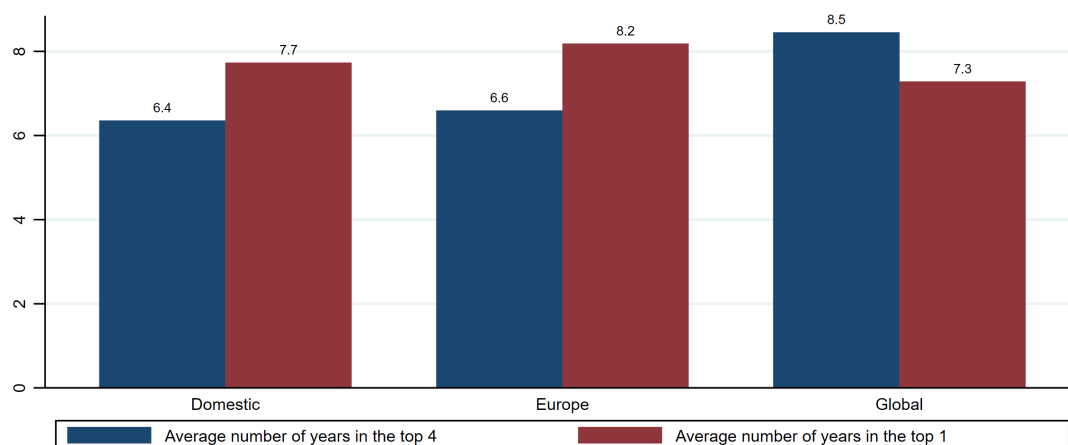
The high share of markets in which there is at most one new firm entering every year reflects low dynamism at the top. This idea is reinforced by looking at Figure 2.16, which shows the average number of years for which a firm remains in the top, either in the top 4 or as a market leader (top 1). Overall, the number of years that a firm remains at the top is high, between 6 and 8 years, suggesting low contestability among the top firms. In industries that compete at the European and the domestic level, on average, a firm remains more as a market leader than in the top 4 (European: 6.6 in the top 4 vs. 8.2 years in the top 1; Domestic: 6.4 in the top 4 vs. 7.7 years in the top 1), suggesting that these industries might be characterised by a dominant firm, which is rarely challenged as a leader, and exhibit relatively more churning among the other firms in the top 4. Conversely, for industries competing at the global level, the average number of years that a firm remains in the top 4 is higher than the average number of years that a firm remains as leader, 8.5 vs 7.3 years, respectively. Overall, these results suggest that industries competing

firms. The preference accorded to the entrenchment measure in this report derives from the fact that – as explained Section 3 - it is built using data from Orbis. As it is well known (Bajgar, Berlingieri, et al. (2020)), Orbis has limited coverage of the population of firms (especially small and medium enterprises) and significant coverage differences over time. Therefore, Orbis is not well suited for an analysis that needs to account for the entry and exit of firms. Note that the two measures are equivalent in a dataset where the firms' population is present.

⁴⁶These categories correspond to markets where each year, with respect to the previous one, there are, respectively, zero, one, two, three, or four new firms in the top 4 group.

at the global level might be characterised by big firms that are persistently among the top 4 and are competing with each other to be the largest firm, while industries competing at the European and the domestic levels have more competition for the top 4 but less competition for the market leadership.

Figure 2.16: Number of years that a firm stays in the top 4 and top1



Note: the chart shows the average number of years that a firm remains in the top 4 and top 1 across industries (and countries, for the domestic bucket) between 2000 and 2019. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are classified as either domestic, European, or global, depending on the results of the taxonomy. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE for the domestic and European bucket, while in the global one also JAP, KOR, and USA are present.

Source: OECD calculations.

Rank persistence

Entrenchment measures the persistence of firms in the group of the top 4, but it does not consider whether there is competition and dynamics among the top 4 firms. This subsection further deepens the latter point, exploring the changes in the ranking of the firms at the top.

Even if there are few changes among the firms that are in the top 4, a market could be highly contestable if there is high competition among them and the top 4 firms constantly challenge each other. Additional evidence on this point is reported in Figure 2.27, 2.28, and in Figure 2.29 in Section 2.B.1 of the Appendix. In these figures, the focus is, respectively, on industries competing at the domestic level that do not have any entry in the top 4, on industries competing at the domestic level with one new firm entering the top 4, and on industries competing at the European level with no new firms in the top 4. These combinations represent the industries belonging to the blue and red bins in the top left graph of Figure 2.15, and the red bin from the top right one⁴⁷. In this exercise, there is a zoom-in in each of these bins to understand, given that there has been 0 or 1 entry in the top 4 group, how many ranking changes there have been among the top 4.

⁴⁷Among all the possible markets and number of new firms in the top 4 (see Figure 2.15), this exercise reports results only for these three categories not only because they represent most of the markets, but also because they are those for which it has been possible to uncover stable and meaningful patterns. The others – available upon request – do not provide any interesting information about rank persistence at the top.

Figure 2.27 shows that for industries competing at the domestic level in which there has not been any entry of new firms in the top 4, most of the markets have no changes in the ranking position of the top 4 firms, about 95 over 150 markets. If there are changes in the top 4, in most of the cases, they just involve two firms that switch their position (about 45 markets). Figure 2.28 focuses instead on industries competing at the domestic level, with one new firm entering at the top 4. Most of those markets have just one change in the top 4, about 60 over 150 markets in 2000, meaning that the new firm in the top 4 is replacing another firm that exits from the top 4, but the other firms remain in the same position. On the other hand, about 70 markets have either two or three changes, suggesting that the entry of one new firm in the top 4 is accompanied by some reshuffling among firms that remain in the top 4. More importantly, the trend with just one change is increasing and diverging from the other lines, hinting that one entry in the top 4 not accompanied by reshuffling at the top is happening more in more recent years than at the beginning of the period. For industries competing at the European level, Figure 2.29 shows that when no new firms enter the top 4 group, most of the industries also have no changes in the ranking of the top 4 firms. Furthermore, the increasing trend of this line compared to the others reported in the same figure suggests that, increasingly often, not only is there no change in the top 4 group composition, but the ranking of the firms remains constant (suggesting low ranking contestability). Overall, the evidence reported so far on entrenchment and rank persistency suggests some lack of dynamism at the top, both with few firms contending the top 4 positions and low levels of contestability among the market leaders.

As a last exercise to investigate market dynamism among the top firms, a transition matrix of the ranking between $t - 1$ and t is studied. Table 2.3 shows the probability of transitioning from a given position in $t - 1$ (one of the different rows in the first column) to a different ranking in t (one of the different columns in the first row). The probabilities are computed by pooling together all geographic markets and years. Table 2.3, once more, suggests rather high persistency among market leaders. Firms that are in a certain position are more likely to remain in their position than to switch to any other one (the diagonal of the matrix). For market leaders, the likelihood of remaining leader, somehow consistently with the high values of leadership ratio reported in Figure 2.11 and the previous evidence from Figure 2.16, is very high (above 80%). As expected, there is lower persistence in the lower rankings (the probabilities are more evenly distributed across the different columns) because firms that exit the top 4 (last column) and firms that enter the top 4 (last row) are more likely to do it at lower ranking. For example, firms that enter the top 4 will do so in 50% of the cases in the fourth position, 24% of the cases in the third position, 15% in the second position and 10% of cases in the first position. Finally, the probability of a firm going to a lower ranking is higher than going up⁴⁸.

2.6 Conclusions

This chapter is part of a larger project conducted at the OECD, and contributes along several dimensions to the growing literature documenting recent trends on market concentration, markups, entrenchment, and other measures that can proxy trends in competition.

⁴⁸In unreported robustness check, available upon request, the transition matrix has also been computed by geographical buckets separately. The results for the transition matrices of each geographical bucket are qualitatively similar to the results from the transition matrix with all the geographical buckets aggregated and suggest the same underlying patterns in the ranking movements in the top 4. As an additional check, the transition matrix has been computed also looking at rank changes using a different time horizon, such as between $t - 3$ and t . While the probabilities are, of course, different and the likelihood of exiting the top 4 (last column) is now much higher, the relative probabilities of transitioning within the top 4 are qualitatively similar to Table 2.3.

Table 2.3: Transition matrix

lt_t	Ranking = 1	Ranking = 2	Ranking = 3	Ranking = 4	ext	prob_0
Ranking = 1	0.806999981	0.093999997	0.021	0.007	0.071000002	1
Ranking = 2	0.093000002	0.635999978	0.137999997	0.035999998	0.097000003	1
Ranking = 3	0.02	0.128999993	0.519999981	0.165999994	0.165000007	1
Ranking = 4	0.008	0.032000002	0.151999995	0.430000007	0.377999991	1
entry	0.104000002	0.155000001	0.237000003	0.503000021		1

Note: the table shows the probability of a firm that is a given ranking position at t-1 to be in a certain ranking position 1 year later, between 2000 and 2019. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE for the domestic and European bucket, while in the global one also JAP, KOR, and USA are present. The probabilities are calculated aggregating all geographical buckets and years.

Source: SBS, OECD own calculations.

It covers in an integrated way multiple measures that proxy for different facets of competition (or a lack of it): concentration and entrenchment. Examining different characteristics of markets in the same setting provides richer inferences about the trends in market power, as well as possible explanations for these trends. The analysis is conducted for both manufacturing and services (and beyond, as it also includes mining and utilities sectors) across many OECD countries and over a long-time horizon. Constructing a database that allows such analysis is itself a significant contribution.

The project innovates on the existing literature in its measurement of concentration, aimed at reflecting markets more accurately. First, it measures concentration within more narrowly defined industries than most previous cross-country studies, mainly at the 3-digit level. Second, international trade is accounted for by constructing a taxonomy that defines whether markets compete domestically or internationally and computes concentration at the corresponding geographic level. Therefore, industries are classified into three geographical buckets: domestic, European, or global. Using this taxonomy means that the concentration measures account for firms' international activities, even when firm-level trade data is absent. On top of this, in the robustness checks, imports and exports are incorporated in the concentration measure using industry-level data. Third, following the methodology by Bajgar, Berlingieri, et al. (2019), the linkages of firms within business groups are accounted for to incorporate the role of mergers and acquisitions in driving concentration trends and to capture the complete activities of multinational firms in a market.

Alongside concentration, a static measure of market shares at any point in time, the project also measures entrenchment. Entrenchment is a dynamic measure of the persistence of firms as market leaders and provides richer insights into the extent of competition, even when concentration is high. The measurement of entrenchment innovates on previous literature by defining markets tradeable internationally following the taxonomy and by accounting for the connectedness of firms within a business group.

In terms of empirical results, on average, industry concentration has increased across all geographic buckets. Industries that compete at the domestic level had the greatest increase in average concentration, by around 6 percentage points (p.p., henceforth) between 2000 and 2019. Industries that compete internationally – either at the European or global level – increased their concentration by approximately 4 p.p.

The comparison of unweighted and weighted concentration cumulative changes shows that when weighting for the relative importance of the markets, concentration looks overall flat over the period 2000-2019 for industries competing both domestically and at the European level. In industries competing at the global level, the weighted trend is even decreasing. The combination of these results suggests that for the domestic and European buckets, the increase in concentration mostly occurs in relatively small markets (in terms of gross output), while for the global buckets, the decrease occurs in relatively big sectors.

Finally, with respect to entrenchment, which captures the persistence of firms at the top, the trends remained relatively flat across all geographical buckets in the period considered (2000-2019).

2.A Data

2.A.1 Production data

National accounts (NA) The variable of interest extracted from NA is the annual value of gross output expressed in millions of Euros for the years 2000-2019. It is available at the Nace2 A*64 level of aggregation, which comprises 66 industries⁴⁹.

The original NA dataset on which the apportioning procedure is based is composed of the 15 countries of the final sample (Belgium, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Norway, Poland, Portugal, Slovenia, Spain, Sweden, and United Kingdom) and by 54 industries. It includes 16200 observations, out of which 294 (1.81%) are missing values.

Structural Business Statistics (SBS) SBS contains information on economic activity of all economic sectors, excluding agriculture and personal services, and provides data at different levels of aggregation, from 1-digit to 4-digit industry level⁵⁰. The data is generally collected by national statistical institutes by means of statistical surveys, business registers and other administrative sources.

Statistical Classification of Economic Activities in the European Community (NACE): NACE Rev.1 was used until 2001, NACE Rev. 1.1 was used from 2002 to 2007, and NACE Rev. 2 was used from 2008 onwards. Available data are reported according to NACE Rev. 1.1 (pre-2008) or NACE Rev 2 (post-2008). Some data are reported both in NACE Rev.1.1 and NACE Rev. 2 for the year 2008.

The types of variables present in SBS, broadly defined, are business demographic variables, output-related variables, and input-related variables. In particular, the variables used to build the final dataset, all available at the country-industry (3-digit) level, are the following:

- **Production value (millions of Euros)** measures the amount produced by the unit, based on sales, including changes in stocks and the resale of goods and services⁵¹.
- **Turnover (millions of Euros)** comprises the totals invoiced by the observation unit during the reference period and corresponds to market sales of goods or services supplied to third parties.
- **Number of persons employed** is defined as the total number of persons who work in the observation unit, as well as persons who work outside the unit who belong to it and are paid by it.
- **Number of enterprises** is a count of the number of enterprises active during at least a part of the reference period.

Production value is the main variable of interest. The other variables are used to conduct consistency checks or as regressors during the imputation process (described below). In the two next sections are reported various statistics on the two SBS datasets. They all refer to the final set of 15 European countries (Belgium, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Norway, Poland, Portugal, Slovenia, Spain, Sweden, and United Kingdom) and years (2000-2019).

⁴⁹For the extended documentation, see: National Accounts metadata; for further definitions, see: European system of accounts - ESA 2010

⁵⁰For further details please refer to: Methodological manual on European Structural Business Statistics – 2021 edition and the SBS website Structural business statistics

⁵¹In this context, unit refers to the enterprises observed to build the SBS data. In this work, only the final data, aggregated at the industry level (as available on Eurostat website), are used.

SBS NACE Revision 1 (pre-2008) Data on pre-2008 years are reported according to NACE Rev. 1.1 industry classification, and cover sections C (Mining and Quarrying), D (Manufacturing), E (Electricity, gas, and water supply), G (Wholesale and retail), H (Hotels and restaurants), I (Transport, storage, and communication), and K (Real estate, renting and business activities) for the years between 2000 and 2008. At the 3-digit level, there are 181 industries and 24,435 observations, of which 5,903 (24.16%) are missing the production variable. The share of missing values varies across countries and sectors, as shown in Table 2.4 At the 3-digit level, sections H, I, and K have the highest share of missing industries⁵².

Table 2.4: SBS NACE Rev.1.1 – Raw Data

Sector	Non-Missing	Missing	Tot. Obs.
C (Mining and Quarrying)	1,118 (63.70%)	637 (36.30%)	1,755
D (Manufacturing)	12,303 (88.48%)	1,602 (11.52%)	13,905
E (Electricity, gas, water supply)	368 (68.15%)	172 (31.85%)	540
G (Wholesale and Retail)	2,461 (95.95%)	104 (4.05%)	2,565
H (Hotels and Restaurants)	91 (13.48%)	584 (86.52%)	675
I (Transport, Storage, and Comm.)	961 (50.85%)	929 (49.15%)	1,890
K (Real Est., Renting, and B. Act.)	1,230 (39.61%)	1,875 (60.39%)	3,105
Total	18,532 (75.84%)	5,903 (24.16%)	24,435

Note: the table presents statistics on missing and non-missing values for 3-digit industries for the sample of countries included in the analysis (BEL, DNK, FIN, FRA, GBR, GER, GRC, HUN, ITA, NOR, POL, PRT, SVN, SPA, and SWE).

Source: SBS, OECD own calculations.

To provide an idea of the quality of the coverage over time, Table 2.5 reports the number and percentages of missing values within a country-industry pairs: among a total of 2,715 industry-country pairs (15 countries by 181 3-digit industries), 54.51% are complete time-series without any missing observation, 11.68% have only one year missing. For time series interpolation of missing values within a country-industry, the most concerning cases are those with few non-missing observations: 522 (19.23%) industry-countries pairs have less than five years of available data. For these cases, methods other than interpolation (described in the following section) have been adopted to impute some observations.

SBS NACE Revision 2 (post-2008) Data from 2008 onwards follow the NACE Rev. 2 industry classification and cover sections B (Mining and Quarrying), C (Manufacturing), D (Electricity, gas, steam and air conditioning supply), E (Water supply; sewerage, waste management and remediation activities), G (Wholesale and retail trade), H (Transportation and Storage), I (Accommodation and food service activities), J (Information and communication), M (Professional, scientific and technical activities), and N (Administrative and support service activities). At the 3-digit level there are 204 industries and 36,516 observations, of which 3,574 (9.79%) with missing production variable – a much lower percentage than the pre-2008 sample. As in the pre-2008 SBS dataset, coverage varies across countries and sectors, as shown in Table 2.6⁵³. In addition, Table A 2.5 reports the distribution of missing values within a country-industry pairs. With respect to the

⁵²Sorted by percentage of missing values: Greece (45.43%), Denmark (35.17%), Belgium (32.41%), Slovenia (27.75%), Portugal (27.5%), Sweden (26.97%), Norway (24.74%), Poland (23.51%), Finland (23.08%), France (21.73%), UK (20.69%), Hungary (18.85%), Spain (17.62%), Germany (12.95%), and Italy (5.59%).

⁵³Sorted by percentage of missing values: Finland (17.85%), Denmark (17.69%), Slovenia (16.5%), Sweden (13.03%), Belgium (12.99%), Norway (12.91%), Germany (10.87%), Greece (7.43%), France (6.66%), UK (5.53%), Portugal (6.5%), Spain (6.37%), Poland (4.41%), Italy (3.92%), and Hungary (3.8%).

Table 2.5: Number of missing observations within country-sector pairs

Number of missing years	Number of country-sector pairs	Percent of country-sector pairs (%)	Cumulative share (%)
0	1,480	54.51	54.51
1	317	11.68	66.19
2	108	3.98	70.17
3	184	6.78	76.94
4	104	3.83	80.77
5	29	1.07	81.84
6	31	1.14	82.98
7	22	0.81	83.79
8	43	1.58	85.38
9	397	14.62	100
Total	2,715		

Note: the table presents statistics on missing and non-missing values for 3-digit industries for the sample of countries included in the analysis (BEL, DNK, FIN, FRA, GBR, GER, GRC, HUN, ITA, NOR, POL, PRT, SVN, SPA, and SWE).

Source: SBS, OECD own calculations.

pre-2008 dataset, there are much less concerning cases: only 5.82% of the country-industry pairs have less than 5 years available in their time series.

Table 2.6: SBS Nace Rev.2 – Raw Data

Sector	Non-Missing	Missing	Tot. Obs.
C (Manufacturing)	15,282 (89.87%)	1,723 (10.13%)	17,005
D (Electricity, etc.)	438 (81.56%)	99 (18.44%)	537
E (Water supply,waste, etc.)	968 (90.13%)	106 (9.87%)	1,074
G (Wholesale and Retail)	3,740 (99.49%)	19 (0.51%)	3,759
H (Transp. And Storage)	1,950 (72.63%)	735 (27.37%)	2,685
I (Accomod. And Food serv.)	1,235 (98.56%)	18 (1.44%)	1,253
J (Inform. And Comm.)	2,202 (94.63%)	125 (5.37%)	2,327
M (Prof. and tech. activities)	2,6413 (98.36%)	44 (1.64%)	2,685
N (Administrat. And supp. Act.)	3,308 (97.27%)	93 (2.73%)	3,401
Total	90,21 (90.21%)	3,574 (9.79%)	35,516

Note: the table presents statistics on missing and non-missing values for 3-digit industries for the sample of countries included in the analysis (BEL, DNK, FIN, FRA, GBR, GER, GRC, HUN, ITA, NOR, POL, PRT, SVN, SPA, and SWE).

Source: SBS, OECD own calculations.

The imputation procedure for missing SBS data Several imputation procedures are used to fill missing values whenever possible. The aim is to fill the gaps in the data while being as conservative as possible in terms of the imputed values. This section briefly describes the steps used to impute the missing values and reports the number of cases filled in by each of them. The steps used to fill in gaps in the data within each country-industry time series – implemented in sequential order are the following:

1. **Accounting Identities:** first, the hierarchy of industry classifications is used here. When-

Table 2.7: Number of missing observations within country-sector pairs

Number of missing years	Number of country-sector pairs	Percent of country-sector pairs (%)	Cumulative share (%)
0	2,283	74.61	74.61
1	210	6.86	81.47
2	116	3.79	85.26
3	77	2.52	87.78
4	66	2.16	89.93
5	46	1.5	91.44
6	44	1.44	92.88
7	40	1.31	94.18
8	27	0.88	95.07
9	21	0.69	95.75
10	34	1.11	96.86
11	34	1.11	97.97
12	62	2.03	100
Total	3,060	100	

Note: the table presents statistics on missing and non-missing values for 3-digit industries for the sample of countries included in the analysis (BEL, DNK, FIN, FRA, GBR, GER, GRC, HUN, ITA, NOR, POL, PRT, SVN, SPA, and SWE).

Source: SBS, OECD own calculations.

ever there is only 1 missing 3-digit industry within a 2-digit industry, the difference between the industry and the sum of the non-missing sectors can be assigned to the missing industry. Analogously, in the case of a missing 2-digit industry, the value of the sum of the 3-digit sectors can be imputed if they are all non-missing.

- 2. Interpolation:** second, if production is missing in one year, the average between the previous and the following year's value is imputed if they are both non-missing.
- 3. Pseudo-Propensity Score Matching (Pseudo-PSM):** production is missing in t but present in $t - 1$ and $t - 2$ such that the growth rate from $t - 2$ to $t - 1$ can be computed. First, the growth rate in t is proxied using the average growth of that industry in t in the five countries that had the closest growth rate in $t - 1$. Subsequently, the production value in t is imputed by attributing the imputed growth rate in t from the first step to the non-missing value in $t - 1$.
- 4. Regressions:** production is regressed on one of the other available relevant variables (turnover, persons employed, and number of enterprises) within each country-industry. The prediction from the estimated regression coefficients is then used to fill the production missing values. This procedure is carried out separately for each relevant industry-country pair.

Over an initial average of 15.68% missing observations, the imputation procedure adopted allows to fill in 10.7% of observation, leaving a final sample with 4.88% of missing observations⁵⁴. Whenever possible, accounting identities and interpolation, being more conservative than the other steps, have been used to impute missing values. On average, across NACE Rev.1.1 and NACE Rev.2, these two methods account for about 4% of the imputation of missing values. Pseudo-PSM allowed

⁵⁴The weights are given by the relative size of the two samples (pre-2008, 0.41%, and post-2008, 0.59%) in the total SBS sample.

to further fill in 3% of missing observations. This method was chosen among others after a careful assessment in terms of out-of-sample prediction: compared with other regression-type prediction methods, it introduces smaller errors⁵⁵. Finally, regressions, which exploit information from other relevant variables (turnover, persons employed, and number of enterprises), accounted for about an additional 4% of observations. The following two tables show the number of observations to which each imputation method is applied, as well as the number of observations that are originally non-missing and the number that cannot be imputed. Table 2.8 shows the pre-2008 data, while Table 2.9 shows post-2008 data.

Table 2.8: Imputation on pre-2008 SBS

Imp. Method	Frequency	Percentage	Cum.
Original	18,532	75.84	75.84
Still Missing	2,823	11.55	87.4
Pseudo PSM	1,074	4.4	91.79
Regressions	866	3.54	95.33
Hierarchy	698	2.86	98.19
Interpolation	442	1.81	100
Total	24,435	100	

Note: the table presents statistics on missing and non-missing values for 3-digit industries for the sample of countries included in the analysis (BEL, DNK, FIN, FRA, GBR, GER, GRC, HUN, ITA, NOR, POL, PRT, SVN, SPA, and SWE).

Source: SBS, OECD own calculations.

Table 2.9: Imputation on post-2008 SBS sample

Imp. Method	Frequency	Percentage	Cum.
Original	32,939	90.2	90.2
Regressions	1,549	4.24	94.45
Pseudo PSM	878	2.4	96.85
Hierarchy	647	1.77	98.62
Interpolation	407	1.11	99.74
Still Missing	96	0.26	100
Total	36,516	100	

Note: the table presents statistics on missing and non-missing values for 3-digit industries for the sample of countries included in the analysis (BEL, DNK, FIN, FRA, GBR, GER, GRC, HUN, ITA, NOR, POL, PRT, SVN, SPA, and SWE).

Source: SBS, OECD own calculations.

Getting disaggregated and harmonised gross outputs Three main steps are implemented to obtain harmonised time series for gross output that span over a very long period, are expressed in the same industry classification, and at a very disaggregated level of aggregation. The first two steps, which involve the filled SBS data, aim at getting for each country-industry pair a continuous time series expressed in a single classification, i.e., NACE Rev. 2, for the entire period of interest, 2000-2019. In order to do so, first, the pre-2008 sample is converted from NACE Rev. 1.1 to

⁵⁵To compute out-of-sample prediction errors, some original non-missing variables are set to zero, and the error is defined as the difference between original observation artificially removed and the value imputed with the different methods under assessment.

NACE Rev. 2. Second, to increase the quality and consistency of the data pre- and post-2008 time series, the converted to NACE Rev. 2 pre-08 data have not been used directly. Instead, they have been used to compute growth rates in each country-industry, and then these growth rates have been used to fill backwards the NACE Rev. 2 series. Finally, 2-digit level NA data have been apportioned to the corresponding 3-digit industries by using the detailed 3-digit SBS data as weights. These steps are described in more detail in the paragraphs below.

The conversion procedure The conversion of industry classifications from NACE Rev. 1.1 to NACE Rev. 2 follows the correspondence tables available on Eurostat – RAMON, and additional details can be found in the NACE Rev. 2 documentation⁵⁶. Many of the NACE Rev 1.1 industries uniquely map to a single NACE Rev. 2 industry (they are 1-to-1 or n-to-1). However, importantly, there are also many cases where a single or multiple NACE Rev 1.1 industries maps to multiple NACE Rev. 2 industries (1-to-m or n-to-m), and these cases represent the most problematic situations to deal with when trying to classify the whole dataset with NACE Rev. 2 classification. The four possible types of correspondences are:

- 1-to-1 correspondences: 195 classes in NACE Rev. 1.1 correspond exactly to one class in NACE Rev. 2 and vice-versa (38%)
- n-to-1 correspondences: 86 cases, where two or more classes in NACE Rev 1.1 correspond to one class in NACE Rev. 2 (17%)
- 1-to-m correspondences: 18 cases, where one NACE Rev. 1.1 class is split into two or more classes in NACE Rev 2 (3%)
- n-to-m correspondences: 215 cases, where two or more classes in NACE Rev. 1.1 correspond to two or more classes in NACE Rev. 2. (42%)

Filling backward It is not straightforward to obtain a consistent time series of production for each country-industry for the whole timespan (2000-2019): even for the easiest cases of 1-to-1 or n-to-1 correspondences, the values in the overlapping year in the two datasets, 2008, are not coinciding⁵⁷.

To obtain the most consistent time series possible spanning pre- and post-2008, the procedure put in place aims at considering both the discrepancies between the two datasets and the industry conversion issue at the same time. Once pre-2008 data (available only in NACE Rev. 1.1) are converted to NACE Rev. 2, they are then used to compute pre-2008 growth rates for each country-industry. Subsequently, they are then applied backwards to the post-2008 time series to obtain a pre-2008 time series consistent in terms of absolute values with the post-2008, while at the same time keeping into consideration the pre-2008 growth rates. See the box below for further details on the methodology applied and Box A.2 for a practical example of the exercises.

⁵⁶Please refer to Europa - RAMON - Correspondence Tables List; NACE Rev. 2 - Statistical classification of economic activities - Products Manuals and Guidelines - Eurostat (europa.eu)

⁵⁷To provide an illustrative example, NACE Rev. 1.1 industry DA151 is mapped 1-to-1 into industry C101 but the value of production reported for Belgium in SBS NACE Rev. 1.1 is 5508.281 while in SBS NACE Rev. 2 is 5148.8.

Example of 3:2 Correspondence

This approach allows us to overcome challenges related also to all the correspondence cases mentioned before. In the case of n-to-1, where many industries in NACE Rev. 1.1 are allocated to a single industry in NACE Rev. 2, the growth rate of the sum of the n industries is applied backwards to the single industry of NACE Rev. 2. In the case of 1-to-m, where a single industry in NACE Rev. 1.1 is allocated to many industries in NACE Rev. 2, the growth rate of the single industry in NACE Rev. 1.1 is applied backwards to all industries in NACE Rev.2. Finally, in the most complicated cases of m-to-n correspondences, the growth rate of the sum of the many relevant pre-2008 industries is applied to each of the relevant post-2008 industries. Consider, for example, the case of a 3:2 conversion, as in the figure below: the three industries y1, y2, and y3 in NACE Rev. 1.1 are converted into two industries, a1 and a2 in NACE Rev. 2. This approach uses the growth rate of the sum of y1 and y2 for a1, and the growth rate of the sum of y1, y2, and y3 for a2. Note, therefore, that the growth rate of a1 is not the same as the growth rate of a2.




Figure 2.17: Example of correspondence

Numerical example of filling backwards

As an example, consider industry SBS NACE Rev. 1.1 DA151 that maps 1-to-1 into NACE Rev. 2 industry C101. As an example, data of a country of the dataset for the years 2005-2008 are reported in Table 2.18. The series of industry DA151 (NACE Rev. 1.1, “SBS NACE 1” in the table) breaks after 2008, while the series C101 (NACE Rev. 2, “SBS NACE 2” in the table) is available only from 2008 onwards. The growth rate is calculated for the NACE Rev. 1.1 series as $g_t = \frac{y_t^1}{y_{t-1}^1}$ (“Growth SBS NACE 1” in the table) and applied to the NACE Rev. 2 series from 2008 backward: $y_{t-1}^2 = \frac{y_t^2}{g_t}$ (“Final SBS series” in the table).

	2005	2006	2007	2008
SBS NACE 1 (y_t^1)	4795.7	5147.5	4985.98	5508.28
Growth SBS NACE 1 (g_t)		1.073	0.968	1.104
SBS NACE 2 (y_t^2)	.	.	.	5148.8
Final SBS series	4482.72	4811.56	4660.59	5148.8

Figure 2.18: Filling Backward Example

The apportioning procedure The final dataset obtained from SBS data contains a measure of gross output (in millions of euros) at the 3-digit NACE rev.2 level for 204 industries and 15 countries from 2000 to 2019. The final step of the procedure is to use this filled SBS dataset to obtain a disaggregated version of the NA data, otherwise available only at A*64 (i.e., slightly more aggregated than the 2-digit level) classification level. In order to get production data at the 3-digit level, the 3-digit SBS data are used to construct shares of the 3-digit industries within each 2-digit industry, which are then used as weights to distribute each 2-digit production data from the NA to 3-digit industries⁵⁸. The 3-digit production data obtained are therefore consistent with NA – which are often based on the population of firms – and at the same time available at the desired level of granularity. The main reason for using the SBS 3-digit production data to apportion the 2-digit ones coming from NA instead of relying directly on them is the following: SBS captures the structure of the economy at a higher level of disaggregation than NA and data are representative of the economy within countries but not across countries (due to different methodologies adopted by national statistical agencies). Therefore, SBS is a good source for obtaining long time series of 3-digit weights but cannot be used directly as a source for measuring production in a cross-country context, making NA necessary to guarantee cross-country comparability.

A key point in the methodology is that SBS provides information on production both at the 2-digit and the 3-digit level. Each 2-digit sector is composed of one or more 3-digit industries. Therefore, it is possible to attribute the share of production that each of them represents in the corresponding 2-digit sector. Importantly, these weights are computed within SBS data and hence give consistent information. The SBS surveys provide indeed a consistent picture of the within-country economic activity (although, as mentioned, a less reliable one than NA for cross-country comparisons). These 3-digit level shares are then used as weights to apportion 2-digit NA data into 3-digit industries.

Mathematically, let $GO_{sct}^{3d NA}$ be the 2-digit value of gross output for industry s in country c and year t obtained from the NA data. Then let $GO_{sct}^{3d SBS}$ and $GO_{Sct}^{2d SBS}$ be respectively the 3-digit

⁵⁸According to NACE rev.2 classification 1-digit, 2-digit, and 3-digit are called, respectively, sections, divisions, and groups. These two nomenclatures will be used interchangeably.

and 2-digit gross output obtained from SBS. The NA apportioned values at the 3-digit level are calculated as:

$$GO_{sct}^{3d NA} = \frac{GO_{sct}^{3d SBS}}{GO_{Sct}^{2d SBS}} GO_{Sct}^{2d NA} \quad (2.11)$$

Example on the apportioning procedure

The 2-digit sector C17 contains two 3-digit industries: C171 and C172^a. The share of sector of C17 accounted by industry C171 is defined according to the formula (suppressing country-year index for clarity):

$$\theta_{C171} = \frac{SBS_{C171}^3}{SBS_{C171}^3 + SBS_{C172}^3}, \quad (2.12)$$

and therefore the 3-digit National Account output is given by: $NAC_{1713} = \theta_{C171} NA_{C17}^2$. Table 2.19 reports the numerical example where the definition is applied to obtain 3-digit National Account output.

Industry	C17	C171	C172
Production (SBS)	16756	6166.7	10589.3
Share (θ)	-	0.3681	0.6319
Production NA (NA)	15941	-	-
3-digit NA	-	5867	10074

Figure 2.19: Apportioning procedure example

^aThese industries are, respectively, Manufacture of paper and paper products (C17), Manufacture of pulp, paper and paperboard (C171), and Manufacture of articles of paper and paperboard (C172).

Summary of final sample for production Given the number of initial missing values in some countries and industries, the final sample has been restricted to 15 countries: Belgium, Bulgaria, Germany, Denmark, Spain, Finland, France, Greece, Hungary, Italy, Norway, Poland, Portugal, Slovenia, Sweden, and the UK. It includes 204 industries over 20 years (2000-2019). There are 61,200 observations, out of which 3,885 (6.35%) are missing. The missing values are concentrated in country-industries combinations where either the whole series is missing or the first ten years are missing (those coming from SBS NACE Rev. 1.1). For a complete list of the 151 industries belonging to the final sample of production and their respective geographical buckets of belonging, please refer to Table 2.10 in Annex B. Due to reasons of compatibility with trade data and lack of data in certain 3-digit industries in SBS, some industries in the final production sample have different levels of aggregations: out of the 151 industries, 133 (88%) are 3-digit, 13 (9%) are 2-digit, and 5 (3%) are 1-digit or above.

2.A.2 Firm level financial data

This section provides a summary of the cleaning procedures applied to the Orbis dataset used in the report. Orbis data have been used for two main purposes: i) identifying the top firms in a market in order to compute concentration, leadership ratio and entrenchment measures, and ii)

for markup estimation. While the initial database is the same, some different cleaning procedures have been adopted for the two purposes. Both are described in this Annex.

As explained in Section 3, financial information within Orbis is available both at the business group-level (consolidated financial data aggregated across subsidiaries belonging to the same owner) and at the individual firm-level (unconsolidated information referring to an individual firm). In this report, unconsolidated accounts are used in the analysis⁵⁹. Consolidated accounts are used only in two cases in which they can be considered equivalent to unconsolidated accounts: for independent firms (i.e., firms that are not part of a business group) and for subsidiaries at the bottom of the ownership hierarchy (subsidiaries not owning further subsidiaries) that do not have unconsolidated accounts, because for such firms consolidated and unconsolidated accounts coincide.

The following two sub-sections provide additional details on the data cleaning preparation for the two samples.

Sample for concentration and entrenchment The methodology outlined in Section 4 to compute concentration and entrenchment requires good coverage of both business-group and individual subsidiaries' financial information. To ensure that all economic activity of each group's subsidiaries is captured, information for firms of all sizes and in all sectors are used⁶⁰.

Following Bajgar, Berlingieri, et al. (2023), some steps are taken to improve the coverage of the data (see Bajgar, Berlingieri, et al. (2019) for a discussion on the impact of these steps on the sample). First, the coverage of Orbis has been expanded by using available consolidated information to infer missing years in the unconsolidated information of the same firm and vice versa. Second, to increase the coverage of consolidated accounts of listed firms, the Worldscope database has been used. Worldscope is merged with Orbis through firms' International Securities Identification Number (ISIN) numbers, which uniquely identify listed firms. For some countries, such as the US, Worldscope can improve the coverage of Orbis substantially. The same cleaning rules used for Orbis data have been applied to Worldscope data. Worldscope reports consolidated financial data and contains very similar values to Orbis consolidated data for observations present in both datasets⁶¹.

As discussed in Section 4, the methodology developed to aggregate sales across all subsidiaries operating in a given market only uses unconsolidated sales of each firm. Consolidated accounts are used only to correct the unconsolidated information in cases where the total subsidiary sales exceed group sales (presumably due to inter-company transactions) or where unconsolidated data are missing. In the latter case, if a headquarter company always reports consolidated accounts but unconsolidated accounts only in some years, the missing years in the unconsolidated accounts are interpolated using growth rates of the consolidated accounts and assuming a constant share of unconsolidated accounts relative to consolidated accounts.

⁵⁹As explained in Section 4, the measure of concentration and entrenchment built in this report looks at business group activities rather than at single firms. The approach adopted in this study fundamentally relies upon unconsolidated data of the individual subsidiaries within a business group to identify the precise industry and location of all the subsidiaries belonging to a group and correctly apportion the group sales to the markets in which the business group is active.

⁶⁰Note that this sample is used at an initial stage of the data construction, in the attempt to consider the worldwide sales of business groups across all sectors in which they are active. This is important because it allows to have similar numbers when comparing the consolidated accounts of the headquarters and the sum of the unconsolidated sales of all their subsidiaries. Subsequently, as explained in Section 3, due mainly to data coverage and comparability across countries, the sample of countries and industries is restricted to those specified in the report.

⁶¹See Bajgar, Berlingieri, et al. (2019) for a discussion on the comparability between the two sources for firms present in both databases.

2.A.3 Ownership data

As explained in Section 4, the business group structure is used to apportion the overall sales of the group across all the relevant markets where it is active. To do so, it requires detailed ownership information on parent-subsidiary linkages. The primary source of firm ownership information is Orbis, which is supplemented with data from the Zephyr database of Mergers and Acquisitions (M&As). Importantly, both datasets are provided by Moody's and share a common firm identifier that allows merging the two datasets.

Orbis contains comprehensive information on ownership linkages among firms, extensively used in the existing literature (Cravino and Levchenko (2017), Fons-Rosen et al. (2021)), which allows to detail ownership linkages between shareholders and their subsidiaries, as well as the identity of the global ultimate owner of subsidiaries (calculated at each calendar year from 2007 until 2020). The global ultimate owner is defined as the firm owning at least 50.01% of the total shares of a subsidiary. This is a commonly used threshold for the definition of control of another firm and, thus, to understand whether the subsidiary's financial information is consolidated into the parent accounts.

To calculate ultimate owners, Orbis uses the tree of ownership linkages for each firm and year. They identify each for each firm its shareholder (the immediate owner), then the shareholder's shareholders and so on. So, for each firm, they start at the bottom and work up the tree of ownership linkages until they find a shareholder that is independent (not controlled by anyone) or controlled by an individual. That shareholder is classified as the ultimate owner of the subsidiary firm at the bottom of the tree.

However, in Orbis, the data primarily starts in 2007 and sometimes later for some firms. Thus, the main data source is complemented with the use of the Zephyr M&A database to measure earlier changes in ownership, enabling the construction of a series starting as early as 2000 whenever data allow, as discussed in the following sub-sections⁶². Zephyr database contains deal-level information on M&As from 1997 onwards for European firms, from 2000 onwards for North American firms, and for other geographic areas from 2003. Overall, Zephyr contains about 2 million M&A deals from 2000 to 2020.

In the following sub-sections, a summary of the methodology is provided. For further details and a more complete discussion, please refer to Bajgar, Berlingieri, et al. (2019), Bajgar, Berlingieri, et al. (2023).

Identifying business groups The procedure to clean and harmonise Orbis and Zephyr relies on the work of Bajgar, Berlingieri, et al. (2023). Following their approach, several steps are undertaken to expand the coverage of the ultimate owner from Orbis. The first step is to use Zephyr to identify changes in immediate (rather than global ultimate) owners not available from Orbis. For each deal, Zephyr contains information on the target, acquirer and vendor firms. About 700,000 deals represent either changes in majority ownership – such as a firm increasing from 10% to 51% equity ownership – or a majority owner further increasing its stake – such as a firm increasing from 51% to 60% ownership. Both types of deals allow to identify the immediate owner of each

⁶²Whilst ultimate ownership data starts in 2007, for some firms it is not available until later years. Common approaches to correcting for this in the literature are either to assume that firms without an Orbis ultimate owner are independent or to take data from a recent year - assuming ownership has not changed over time. Both approaches are problematic. With increasing coverage of ownership over time in Orbis, the former approach will falsely equate missing data with independence and lead to an overstatement of ownership changes over time. The latter approach will lead to an understatement of ownership changes over time and will typically overstate the number of markets and countries in which a firm operates.

target firm at the time of the deal. Furthermore, for changes in majority ownership – when the target firm switches hands – the vendor firm represents the previous immediate owner.

A second step is to use the information available from the table “current” Orbis linkages, which provides direct and indirect ownership linkages. These are used to retrieve the identity of the ultimate owner in cases where the latter is missing but it is possible to identify a shareholder with an indirect share higher than 50.01

The third step is to translate the changes in immediate owners (from the first two steps above) to changes in the ultimate owner. The immediate owner who acquired the target firm may not be the ultimate owner. To find the ultimate owner, the same procedure used by Orbis is followed. Zephyr immediate owner and available information on ownership linkages are combined to find the shareholders of the immediate owners, the shareholders of their shareholders, and so on. The 50.01% criterium is used until the procedure arrives to a shareholder that is either independent or controlled by an individual. This final shareholder is deemed the ultimate owner.

The fourth step is to impute missing years of ownership information and information and roll the owner backwards and forwards until there is a M&A or change in ownership (from the steps above). The additional information on ownership changes allows to roll the ownership information forwards and/or backwards until there was a change in owner, rather than simply assuming that a missing ultimate owner implies independence between firms. For example, if firm C is the ultimate owner of firm A in 2010, and from Zephyr M&A data it is known that firm A was acquired in 2008, then the ultimate ownership information is rolled backwards until 2008. Moreover, in about half of the acquisition cases in the M&A sample, it is also known that firm A was acquired from vendor firm B in 2008, so it is possible to infer that firm B was the (immediate) owner of firm A, and roll back further until an earlier M&A transaction.

Data cleaning Numerous steps are undertaken to identify and correct potential issues in the ownership data, especially to identify missing linkages amongst the largest firms. Spot-checking revealed that some large firm groups have missing ownership linkages between the parent firm and their subsidiaries for some years. This can be problematic because it can lead to double counting of group activity, with both the parent’s consolidated financials and their subsidiary information included as separate groups.

Accordingly, the following checks are undertaken to mitigate this risk. First, ultimate owners who are themselves majority-owned by another firm cannot be true ultimate owners and are therefore adjusted in the data. Second, temporary (one or two year) deviations in ultimate owner relationships, whereby a firm’s ultimate owner changes for just one or two years and then reverts to its previous owner, are removed, as this is an unusual phenomenon in ownership and is most likely to be measurement error. These two steps affect approximately 10,000 firms per year.

Third, to detect missing linkages, large firms that change from having no subsidiaries to a large number of subsidiaries from one year to the next are examined and manually updated, if necessary. Spot-checking revealed cases of intermediate holding companies (that often have no financial information) being temporarily incorrectly identified as the ultimate owner. To address this issue, large groups of subsidiaries (in terms of sales) that have a parent with no financials but switch to a new parent in the following period that does have financials are examined manually. Cases of M&As identified by Zephyr have been excluded, and only cases where more than 90% of subsidiaries transfer to the new parent have been considered. The 1500 largest groups identified, corresponding to groups with sales larger than 10 million Euro, have been corrected in the following way. For the 150 largest groups, each group has been manually inspected against their financial statements, while

for other groups, a name-matching algorithm has been used to semi-automate the identification of whether the prior owner was, in fact, a holding company of the new parent. Those with very similar names have been considered as part of the same group, correcting 147 groups.

Fourth, large firms that never have any subsidiary and, vice versa, large groups of subsidiaries that never have a parent with financials are examined to identify missing links. This builds on the previous step, identifying large groups of subsidiaries that never have a parent with financials, and large parents that never have subsidiaries. In total, 1,031 parents with sales of more than 1 billion Euros that never have subsidiaries have been found, and 251 groups of subsidiaries with more than 1 billion Euros of sales that never have a parent with financials. Again, a name-matching algorithm has been used to semi-automate the identification of whether the prior owner was, in fact, a holding company of the new parent. This process applies to cases of groups (large groups of subsidiaries or large parents) with sales larger than 40 million and treats those with very similar names as part of the same group. In total, based on visual inspection of the name-matching string-matching similarity, 287 groups per year have been corrected.

Fifth, missing links where there are ownership changes amongst firms with very similar names – and are so very likely part of the same group (e.g., ABC Motors acquired by ABC Motors Thailand) – are identified and corrected. This considers any ownership change where the owners have a similar root to their name (e.g., “XYZ Inc” and “XYZ Plc”). These remaining firms are not large or do not have completely missing subsidiaries; if they had, they would have been encompassed in the earlier cleaning steps. Therefore, these firms are somewhat less problematic for the resulting concentration metrics. Given this reduced risk and the fact that all firms in the data are considered part of this step, an automated check using name matching is carried out, requiring an identical match of the cleaned name. Common company type abbreviations (e.g., Plc, Ltd, SA, GmbH, etc.), country names (e.g., ABC (Viet Nam) Ltd) and punctuation are removed, and the resulting root of the name is required to be identical. The global ultimate owner is modified only when the ownership change involves two companies with almost exact names, and the ownership change happens between one ultimate owner that has financials and the other one that has no financials. In total, approximately 5,000 cases are corrected.

Finally, a further check for groups with total gross output (considered as the sum of the unconsolidated accounts of all its subsidiaries) in a given country, industry, and year larger than 150 million Euro is conducted. Within this subset of business groups, firms with similar names in the same country and industry are checked using a string-matching algorithm to select relevant cases. As this algorithm captures situations where the ownership tree is partially missing some links, the spotted cases might be particularly relevant for concentration measures. For example, the company ACCO Brands Europe was reported as a GUO, while it is part of the group Acco Brands Corporation. Both were active in the same country-industry and year, therefore leading to a downward biased measure of concentration. In this situation, there are two groups (or simple firms), both with large revenues but that do not have a complete ownership structure. The subsample of GUOs with similar names and active in the same year has been manually inspected in order to understand when the GUO was indeed the same, correcting approximately 300 cases.

2.A.4 Sample for concentration and entrenchment

The final sample used for studying concentration measures and entrenchment spans 20 years (2000-2019) and is composed by 127 distinct industries allocated to the three different geographical buckets (27 are domestic, 80 are European, and 20 are global). Out of these, 112 (88%) are 3-digit, 10 (8%) are 2-digit, and 5 (4%) are aggregation of two or more 2-digit. The difference from the number of industries included in the production sample is due essentially to combined data

limitation either in ORBIS at the firm level, in gross in gross output at the industry level or in the trade data⁶³. The analysis for Concentration and Entrenchment uses all these different data, and so it has been necessary to restrict the focus on a common and harmonized sample across this different information.

⁶³The industries excluded from the concentration sample with respect to those listed in Table 2.10 are 051, 052, 06, 07, 120, 253, 268, 304, 351, 495, 691, 692, 70, 741, 742, 743, 749, 750, 811, 813, 821, 822, 823, 829 (see Table 2.10 of Section 2.B of the Appendix for the sectors descriptions).

2.B Methodology

Industry	Description	Taxonomy	Industry	Description	Taxonomy
051	Hard coal	Global	275	Domestic appliances	European
052	Lignite	Domestic	279	Other electrical equip.	European
06	Petroleum extract.	Global	281	GP machinery	European
07	Mining	Global	282	Other GP machinery	European
081	Stone, sand, clay	European	283	Agriculture machinery	European
089	Mining, quarrying	Global	284	Machine tools	European
091	Petroleum	Domestic	289	Special-purpose machinery	European
099	Other mining	Domestic	291	Motor vehicles	European
101	Meat processing	Domestic	292	Motor bodies	European
102	Fish processing	European	293	Motor parts	European
103	Fruit processing	European	301	Ship building	Global
104	Oils	European	302	Trains	European
105	Dairy	Domestic	303	Air and spacecraft	Global
106	Grains	European	304	Military vehicles	Domestic
107	Bakery	Domestic	309	Transport equip.	European
108	Other food	European	310	Furniture	European
109	Animal feeds	Domestic	321	Jewellery	Global
110	Beverages	Domestic	322	Musical instruments	Global
120	Tobacco	Domestic	323	Sports goods	European
131	Textile fibres	European	324	Toys	European
132	Textile weaving	European	325	Medical instruments	Global
133	Textile finishing	Domestic	329	Other manufacturing	European
139	Other textiles	European	331	Machinery repair	Domestic
14	Animal production	Global	332	Industrial machinery	Domestic
151	Tanning leather	Global	351	Power generation	Domestic
152	Footwear	European	352	Gas	Domestic
161	Sawmilling wood	European	353	Air conditioning	Domestic
162	Wood products	European	360	Water	Domestic
171	Paper	European	37T39	Sewerage; waste collection	Domestic
172	Paper articles	European	45T47	Wholesale and Retail	European
181	Printing activities	Domestic	491	Passenger rail	Domestic
182	Recorded media	European	492	Freight rail	European
19	Coke / petroleum	European	493	Other passenger land	Domestic
201	Basic chemicals	European	494	Road freight	European
202	Pesticides	European	495	Pipeline transport	Global
203	Paints	European	501	Passenger sea	Domestic
204	Detergents	European	502	Freight sea	Global
205	Chemical products	European	503	Passenger inland water	Domestic
206	Man-made fibres	European	504	Freight inland water	European
211	Basic pharmaceutical	Global	51	Hard coal	Global
212	Pharmaceutical	Global	52	Lignite	European
221	Rubber products	European	53	Postal	European
222	Plastic products	European	1 (55.56)	Accommodation; food and beverage ser.	Domestic
231	Glass	European	581	Book publishing	Domestic
232	Refractory products	European	582	Software publishing	Global
233	Clay materials	European	59T60	TV programme and broadcasting act.	Domestic
234	Other porcelain	European	61	Crude petroleum	Domestic
235	Cement, plaster	European	62T63	Comp. programming & Consultancy	European
236	Concrete products	Domestic	691	Legal activities	Domestic
237	Stone cutting	Global	692	Accounting	Domestic
239	Abrasive products	European	70	Head office activities	European
241	Basic iron/steel	European	71	Iron ores	European
242	Steel tubes	European	72	Non-ferrous metal	European
243	Steel first processing	European	73	Advertising	European
244	Precious metals	European	741	Specialised design	Domestic
245	Metal casting	European	742	Photography	European
251	Structural metal	European	743	Translation	Domestic
252	Metal tanks	European	749	Other professional	Global
253	Steam generators	Global	750	Veterinary	Domestic
254	Weapons	Global	77	Rental	European
255	Metal forging	Domestic	781	Employment agencies	European
256	Metal treatment	Domestic	782	Temp agencies	European
257	Tools	European	783	Other HR	European
259	Other fabricated metal	European	791	Travel agency	European
261	Electronic components	Global	799	Other travel	European
262	Computers	European	801	Private security	European
263	Communication equip.	Global	802	Security systems	European
264	Consumer electronics	European	803	Investigation	European
265	Measurement instruments	Global	811	Facilities support	Domestic
266	Electromedical	Global	812	Cleaning	Domestic
267	Optical equip.	Global	813	Landscaping	Domestic
268	Optical media	European	821	Office admin	European
271	Electric motors	European	822	Call centres	Domestic
272	Batteries	European	823	Conventions	Domestic
273	Wiring	European	829	Business services	European
274	Lighting	European			

Industries in red are excluded from the sample for concentration, leadership ratio, and entrenchment.

Table 2.10: Production and concentration sample

2.B.1 Trends

Figure 2.20: Concentration levels across geographical buckets (weighted)

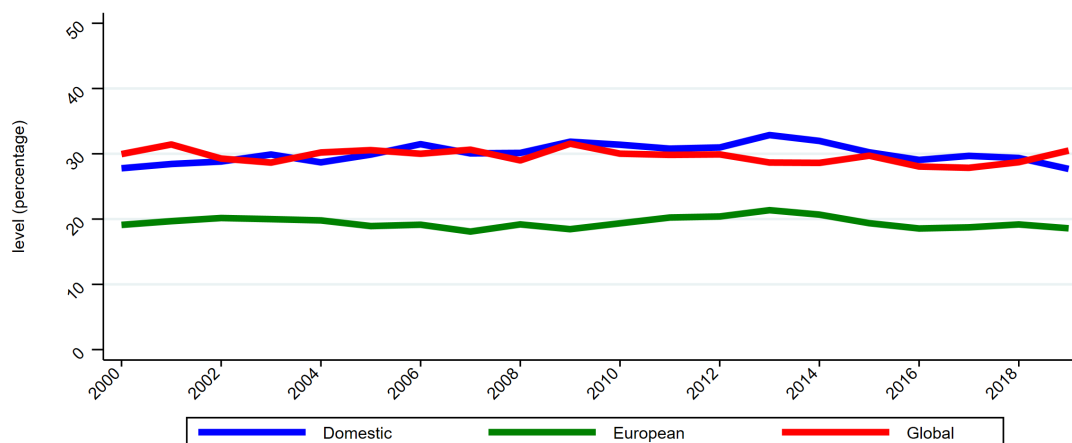




Figure 2.21: Concentration cumulative changes across geographical buckets (weighted)



Note: the chart shows the weighted average across industries (and countries, for the domestic bucket) of CR4 cumulative growth. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are classified as either domestic, European, or global, depending on the results of the taxonomy. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE for the domestic and European bucket, while in the global one also JAP, KOR, and USA are present.

Source: OECD calculations.


Figure 2.22: Concentration level treating all industries as European (unweighted)



Note: the chart shows the weighted average across industries (and countries, for the domestic bucket) of CR4 levels. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are all classified as European. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.


Figure 2.23: Concentration cumulative changes treating all industries as European(unweighted)



Note: the chart shows the weighted average across industries (and countries, for the domestic bucket) of CR4 cumulative growth. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are all classified as European. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.

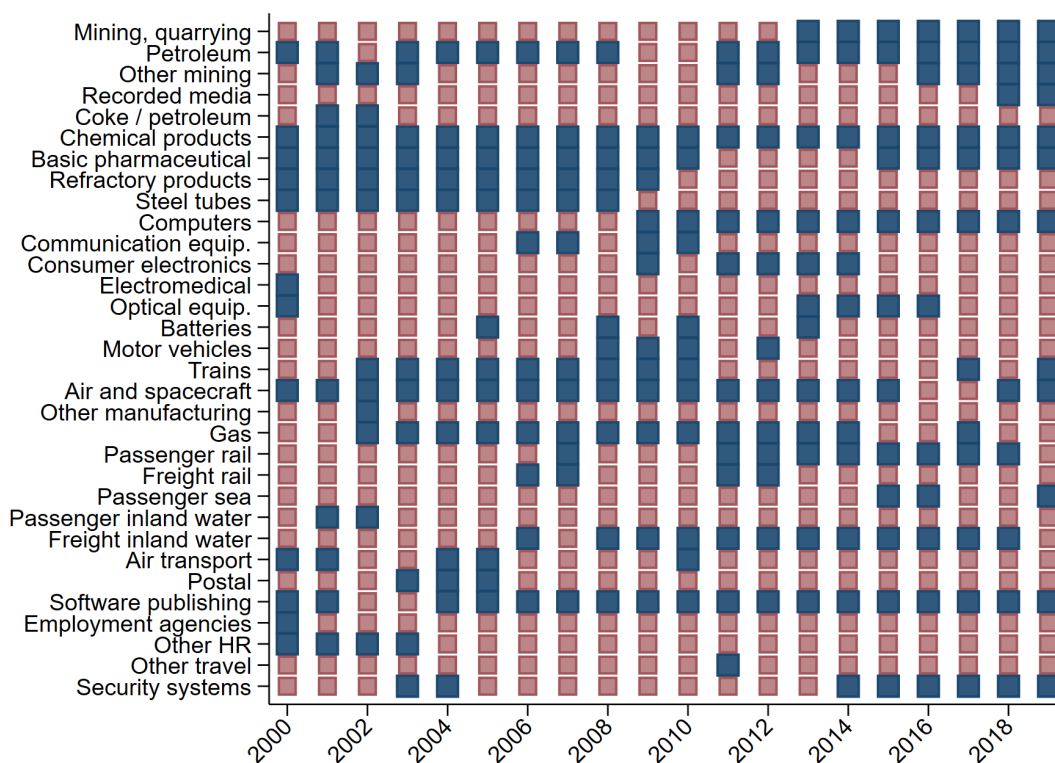
Figure 2.24: Number of distinct industries per decile of concentration



Note: the chart shows the number of industries appearing at least once in each decile throughout the years (each decile in a given year is made of 12 industries). This means, for example, that 32 distinct industries entered the 10th decile (compare Figure A C.6) across the time span of the sample, while 22 appeared at least once in the 1st decile. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are all classified as European. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.

Figure 2.25: Industries in the top decile of the concentration distribution



Note: the chart shows the top decile composition in terms of industries. A given industry can belong to the top decile (blue square) or not (red square) each year. The rows show the presence year by year of each industry (marked by a blue square). The columns show the decile composition in each year. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market services and are all classified as European. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.

Figure 2.26: Leadership ratio across geographical buckets (1st vs. 2nd), weighted

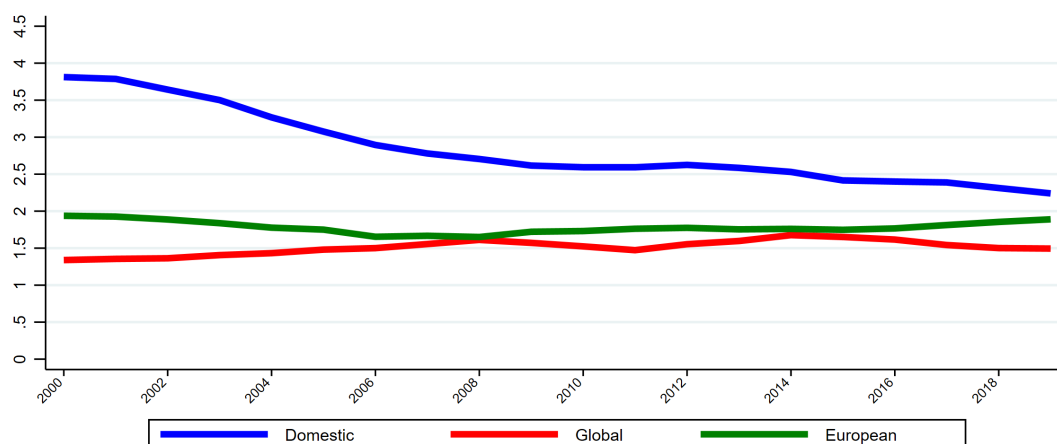




Figure 2.27: Changes in the top4 ranking with no new firms in the top4, domestic bucket



Note: the chart represents the number of changes in the ranking for industries competing domestically in which there is no new firm entering the top4 group. Each line represents the number of industries in which there where, respectively, zero, two, three, and four changes in the ranking of the top 4 firms. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market belonging to the domestic bucket. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.


Figure 2.28: Changes in the top 4 ranking with one new firm in the top 4, domestic bucket



Note: the chart represents the number of changes in the ranking for industries competing domestically in which there is one new firm entering the top4 group. Each line represents the number of industries in which there where, respectively, one, two, three, and four changes in the ranking of the top 4 firms. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market belonging to the domestic bucket. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.

Figure 2.29: Changes in the top4 ranking in markets with no entry, European bucket



Note: the chart represents the number of changes in the ranking for industries competing at the European level in which there is no new firm entering the top4 group. Each line represents the number of industries in which there were, respectively, zero, two, three, and four changes in the ranking of the top 4 firms. Industries included in the analysis are mix of 2 and 3-digit industries belonging to mining, manufacturing, utilities, and non-financial market belonging to the European bucket. The countries included in the sample are BEL, DEU, DNK, ESP, FIN, FRA, GBR, GRC, HUN, ITA, NOR, POL, PRT, SVN, and SWE.

Source: OECD calculations.

References

- Affeldt, Pauline et al. (2021). “Market Concentration in Europe: Evidence from Antitrust Markets”.
- Aghion, Philippe et al. (Nov. 2019). *A Theory of Falling Growth and Rising Rents*.
- Akcigit, Ufuk and Sina T. Ates (2021). “Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory”. In: *American Economic Journal: Macroeconomics* 13 (1).
- Altomonte, Carlo et al. (2021). “Business groups as knowledge-based hierarchies of firms”.
- Amiti, Mary and Sebastian Heise (May 2021). “U.S. Market Concentration and Import Competition”.
- Andrews, Dan, Chiara Criscuolo, and Peter N. Gal (2019). “The Best versus the Rest: Divergence across Firms during the Global Productivity Slowdown”. In: *CEP Discussion Paper* (1645).
- Antoniades, Alexis (Mar. 2015). “Heterogeneous Firms, Quality, and Trade”. In: *Journal of International Economics* 95 (2), pp. 263–273.
- Autor, David, David Dorn, et al. (May 2020). *The fall of the labor share and the rise of superstar firms*.
- Autor, David, Christina Patterson, and John Van Reenen (Apr. 2023). “Local and National Concentration Trends in Jobs and Sales: The Role of Structural Transformation”.
- Bajgar, Matej, Giuseppe Berlingieri, et al. (2019). “Industry concentration in Europe and North America”.
- (June 2020). “Coverage and representativeness of Orbis data”.
- (2023). “Industry concentration in Europe and North America”. In: *Industrial and Corporate Change*.
- Bajgar, Matej, Chiara Criscuolo, and Jonathan Timmis (Dec. 2021). “Intangibles and industry concentration: Supersize me”.
- Barkai, Simcha (Oct. 2020). “Declining Labor and Capital Shares”. In: *Journal of Finance* 75 (5), pp. 2421–2463.
- Benkard, C. Lanier, Ali Yurukoglu, and Anthony Lee Zhang (Apr. 2021). *Concentration in Product Markets*.
- Benmelech, Efraim, Nittai K. Bergman, and Hyunseob Kim (Apr. 2022). “Strong Employers and Weak Employees How Does Employer Concentration Affect Wages?” In: *Journal of Human Resources* 57 (SpecialIssue 1), S201–S250.
- Berlingieri, Giuseppe et al. (2020). “Laggard firms, technology diffusion and its structural and policy determinants”. In: *OECD Science, Technology and Industry* (86).
- Bernard, Andrew B. et al. (2012). “The Empirics of Firm Heterogeneity and International Trade”. In: *Annual Review of Economics* 4 (1), pp. 283–313.
- (June 2018). “Global Firms”. In: *Journal of Economic Literature* 56 (2), pp. 565–619.
- Berry, Steven, Martin Gaynor, and Fiona Scott Morton (2019). *Do increasing markups matter? Lessons from empirical industrial organization*.
- Besley, Timothy, Nicola Fontana, and Nicola Limodio (June 2021). “Antitrust Policies and Profitability in Nontradable Sectors”. In: *American Economic Review: Insights* 3 (2), pp. 251–265.
- Bessen, James E. et al. (2020). “Declining Industrial Disruption”. In: *SSRN Electronic Journal*.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen (2013). “Identifying technology spillovers and product market rivalry”. In: *Econometrica* 81 (4), pp. 1347–1393.
- Calligaris, Sara, Chiara Criscuolo, and Luca Marcolin (2018). “Mark-ups in the digital era”. In: *OECD Science, Technology and Innovation Policy Papers* (No. 94).
- Competition and Markets Authority (Apr. 2022). *The State of UK Competition Report*.

- Covarrubias, Matias, German Gutierrez, and Thomas Philippon (2019). "From Good to Bad Concentration? US Industries over the Past 30 Years". In: NBER.
- Cravino, Javier and Andrei A. Levchenko (May 2017). "Multinational Firms and International Business Cycle Transmission*". In: *The Quarterly Journal of Economics* 132 (2), pp. 921–962.
- Crouzet, Nicolas and Janice Eberly (May 2019). *Understanding Weak Capital Investment: the Role of Market Concentration and Intangibles*.
- Davies, Stephen (Mar. 2022). "Competition and Concentration: Charting the Faultlines".
- De Loecker, Jan, Jan Eeckhout, and Simon Mongey (2022). "Quantifying market power and business dynamism in the macroeconomy".
- De Loecker, Jan, Jan Eeckhout, and Gabriel Unger (May 2020). "The Rise of Market Power and the Macroeconomic Implications*". In: *The Quarterly Journal of Economics* 135 (2), pp. 561–644.
- De Ridder, Maarten (2023). "Market Power and Innovation in the Intangible Economy". In: *American Economic Review*.
- Döttling, Robin, German Gutierrez Gallardo, and Thomas Philippon (July 2017). "Is There an Investment Gap in Advanced Economies? If So, Why?" In: *SSRN Electronic Journal*.
- Eaton, Jonathan et al. (Nov. 2016). "Trade and the Global Recession". In: *American Economic Review* 106 (11), pp. 3401–3438.
- Ekerdt, Lorenz K F et al. (2023). *The Rise of Specialized Firms * †*.
- Fons-Rosen, Christian et al. (July 2021). "Quantifying productivity gains from foreign investment". In: *Journal of International Economics* 131, p. 103456.
- Freund, Caroline and Dario Sidhu (Feb. 2017). "Global Competition and the Rise of China".
- Furman, Jason and Peter Orszag (2018). "A Firm-Level Perspective on the Role of Rents in the Rise in Inequality for us". In: ed. by Martin Guzman. Columbia University Press.
- Gal, Peter N (2013). "Measuring total factor productivity at the firm level using OECD-ORBIS". In.
- Gervais, Antoine and J. Bradford Jensen (May 2019). "The tradability of services: Geographic concentration and trade costs". In: *Journal of International Economics* 118, pp. 331–350.
- Grullon, Gustavo, Yelena Larkin, and Roni Michaely (July 2019). "Are US industries becoming more concentrated?" In: *Review of Finance* 23 (4), pp. 697–743.
- Gutierrez, German and Thomas Philippon (2023). "How European markets became free: a study of institutional drift". In: *Journal of the European Economic Association*.
- Honjo, Yuji, Noriyuki Doi, and Yasushi Kudo (June 2018). "The Turnover of Market Leaders in Growing and Declining Industries: Evidence from Japan". In: *Journal of Industry, Competition and Trade* 18 (2), pp. 121–138.
- Hsieh, Chang-Tai and Esteban Rossi-Hansberg (Mar. 2023). "The Industrial Revolution in Services". In: *Journal of Political Economy Macroeconomics* 1 (1), pp. 3–42.
- IMF (2019). "The Rise of Corporate Market Power and Its Macroeconomic Effects". In.
- Kalemli-Ozcan, Sebnem et al. (2023). "How to Construct Nationally Representative Firm Level Data from the Orbis Global Database: New Facts on SMEs and Aggregate Implications for Industry Concentration". In: *American Economic Journal: Macroeconomics*.
- Koltay, Gábor, Szabolcs Lorincz, and Tommaso Valletti (2022). "Concentration and Competition: Evidence from Europe and Implications for Policy".
- Kwon, Spencer Yongwook, Yueran Ma, and Kaspar Zimmermann (2023). "100 Years of Rising Corporate Concentration". In: *SSRN Electronic Journal*.
- Lashkari, Danial, Arthur Bauer, and Jocelyn Boussard (2019). "Information Technology and Returns to Scale".
- Leamer, Edward (1988). "Measures of Openness". In: ed. by Robert E Baldwin. University of Chicago Press, pp. 145–204.

- Liu, Ernest, Atif Mian, and Amir Sufi (2022). “Low Interest Rates, Market Power, and Productivity Growth”. In: *Econometrica* 90 (1).
- Lyons, Bruce, Catherine Matraves, and Peter Moffatt (Feb. 2001). “Industrial Concentration and Market Integration in the European Union”. In: *Economica* 68 (269), pp. 1–26.
- OECD (2021). *Methodologies to Measure Market Competition*.
- Pellegrino, Bruno (2023). “Product Differentiation and Oligopoly: A Network Approach”. In: *SSRN Electronic Journal*.
- Peters, Michael and Conor Walsh (2020). “Population Growth and Firm Dynamics”. In: (December).
- Rinz, Kevin (Apr. 2022). “Labor Market Concentration, Earnings, and Inequality”. In: *Journal of Human Resources* 57 (Special Issue 1), S251–S283.
- Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Nicholas Trachter (May 2021). “Diverging Trends in National and Local Concentration”. In: *NBER Macroeconomics Annual* 35, pp. 115–150.
- Shapiro, Carl (Nov. 2018). “Antitrust in a time of populism”. In: *International Journal of Industrial Organization* 61, pp. 714–748.
- Smith, Dominic A. and Sergio Ocampo (Feb. 2022). “The Evolution of U.S. Retail Concentration”.
- Stansbury, Anna and Lawrence H. Summers (2020). “The declining worker power hypothesis: An explanation for the recent evolution of the american economy”. In: *Brookings Papers on Economic Activity* 2020 (1).
- Syverson, Chad (Aug. 2019). “Macroeconomics and Market Power: Context, Implications, and Open Questions”. In: *Journal of Economic Perspectives* 33 (3), pp. 23–43.
- Werden, Gregory J. and Luke M. Froeb (2018). “Don’t Panic: A Guide to Claims of Increasing Concentration”. In: *SSRN Electronic Journal*.

Chapter 3

A note on technological change and inequality

3.1 Introduction

In this short note, I briefly touch on a few of the several empirical patterns studied in the macroeconomic literature in the last decades: the increase in top income and top-wealth inequality (see Kopczuk (2015) and Saez and Zucman (2016)), the decline in the entrepreneurship rate (Haltiwanger (2022), Kozeniauskas (2018), and Salgado (2019)), and the so-called divergence in the firm distribution (Andrews, Criscuolo, and Gal (2016), Akcigit and Ates (2021)). Please notice that this is a simple exposition of a possible modelling device rather than a full explanation of these phenomena.

3.1.1 Overview

In this note, I claim that the abovementioned patterns can be explained by the same form of skill-biased technological change used by Poschke (2018), who introduced the so-called entrepreneurial skill-biased technological change. In the literature, general technological progress has been modelled as augmenting the productivity of some specific factor of production (or all of them), allowing firms to scale up their production. Skill-biased technological change instead takes mainly two forms: i) the improvement in the returns of specific skills (skilled labour: see, for example Acemoglu (2002) and Katz and Murphy (1992)) and ii) the decrease of the price of certain factors complementary to skilled labour (decrease in the price of capital, see Krusell et al. (2000)). Poschke (2018) combines general technological progress and skill-biased technological progress. On the one hand, general technological progress brings in new technologies that, in principle, are available to all firms. On the other hand, the available technology level interacts with the single agents' ability: only the most skilled and productive can fully take advantage of the latest and more advanced technological improvements. The most able, or those that run the most productive firms, stay closer to the technological frontier; the others can be considered as laggards (on some stylised facts concerning the firm distribution, refer to Andrews, Criscuolo, and Gal (2016), who point out several dimensions where this divergence phenomenon occurs).

While Poschke (2018) used this assumption to explain the increasing divergence in firm size distribution at different stages of economic development, I see it as a tool to explain some trends in

inequality. In particular, these trends are represented in Figure 3.1 and in Figure 3.2. The former shows the increase in the share of income earned, respectively, by the top 10% and top 1% earners in the United States over the period 1980-2020. The latter shows the analogue quantities for wealth, defined as net worth. Many studies have documented these trends, and several explanations have been advanced. With respect to wealth Kaymak and Poschke (2016) finds that the increase in wage inequality is the main responsible for the increase in wealth concentration. Hubmer, Krusell, and Smith. (2021) instead attributes a higher relevance to the change in the structure of income taxation.


In the first part of this chapter, I build a simple model to show the relationship between skill-biased technological change and top-income inequality and conduct a counterfactual experiment that qualitatively replicates the increase in top-income concentration by assuming that technological change can be proxied by Total Factor Productivity. In the model economy, agents decide whether to be workers and receive a wage or to be entrepreneurs and run a firm whose productivity depends both on the general technological level and on the agent's entrepreneurial ability. When technological progress occurs, the entrepreneurs closer to the frontier become increasingly more productive, and income concentration, which for entrepreneurs is given by their profits, increases. From an economic viewpoint, technological progress increases entrepreneurial returns at an increasing rate with their skills. A by-product of the model is that, in general equilibrium, when firms become larger due to the increase of technology, their size increases and, therefore, also their labour demand. This raises the equilibrium wage and, consequently, the outside option from entrepreneurship. Therefore, the mass of entrepreneurs in the economy decreases, in line with the findings by Salgado (2019). The technology specification also implies a Pareto distribution of entrepreneurial incomes (see Gabaix (2016) for a readable introduction to the use of power laws in economics), without assuming a Pareto distribution in their skills, as it was done in the original article by Lucas (1978).

In the second part, I describe a more elaborate model that I plan to use to study the effect of technological change on wealth inequality, but that I haven't fully solved yet. The dynamic model, in principle, rests on the same mechanisms on which the static model is based, but it uses entrepreneurs to generate a level of wealth concentration that is consistent with the data. The idea, if the model is solved, is to study the impact of technological change on wealth concentration at the top when mediated by entrepreneurial skill-biased technological change.

The model has three main ingredients: the first is the presence of heterogeneous agents like in Aiyagari (1994). This is needed in order to generate an endogenous distribution of wealth. The second is the presence of entrepreneurs in the economy, like in Cagetti and Nardi (2006) or Buera and Shin (2013). The presence of entrepreneurs is important to capture the empirical features of wealth distribution. Entrepreneurs, both empirically and theoretically, have different saving motives than workers: they need to circumvent collateral constraints in order to exploit their ventures fully. The third element is the introduction, following Poschke (2018) and analogously with the static model described above, of skill-biased technological change in the entrepreneurs' technology. The main idea is that while general technological progress increases everybody's firm productivity (at least in the model), the more capable entrepreneurs will gain a relatively larger share. This can explain the increase in wealth concentration and the decrease in the number of entrepreneurs: the increase in wages that follows from technological progress drives toward salaried work marginal entrepreneurs. In the dynamic model, again following Poschke (2018), I add a further assumption on the agents' ability: the skill of an agent if she becomes a worker is correlated with the ability she would have had if she became an entrepreneur. This assumption allows me to generate what can be called subsistence entrepreneurs: low-skilled agents who find it


more profitable to be self-employed in very small firms rather than work for a salary. This feature allows me to study the divergence in firm distribution driven by technological change in the context of a quantitative dynamic macroeconomic model.

Figure 3.1: Evolution of pre-tax income shares earned by the top earners




Source: World Income Database

Figure 3.2: Wealth (net worth) shares owned by the richest fraction of the wealth distribution



Source: World Income Database

Figure 3.3: Labor productivity: frontier vs.laggard firms



Source: Andrews, Criscuolo, and Gal (2016)

3.1.2 Literature review

This paper is related to several strands of literature, and to keep this chapter in the form of a short note, I only mention the closest contributions to the present work.

The first strand of literature related to this paper is that of heterogeneous agents in macroeconomic models, which are crucial to discussing distributional matters. On the production side, the literature on heterogeneous firms is broad, and two of the earliest examples are Lucas (1978) and Hopenhayn (1992). In particular, Lucas (1978) characterises the distribution of firm size based on the managers' skill distribution. In his paper, Lucas shows that if the skill's distribution follows a Pareto distribution, then the firm's size distribution is Pareto. In the static model, I build on a version of Lucas's model, augmenting it with exogenous technological progress and with Poschke (2018) assumption on technology to show how technological progress can influence income distribution.

On the consumer side, heterogeneous agents are needed to study income and wealth distribution, and the origins of the use of such models in macroeconomics can be traced back to Bewley (1986), Imrohoroglu (1989), Huggett (1993), and Aiyagari (1994). The main feature of these models is the presence of idiosyncratic shocks that hit the agents who solve a multi-period consumption-savings problem. Agents use their saving to insure against idiosyncratic shocks, and the heterogeneity in shock histories generate an endogenous wealth distribution. An issue with these models is that they do not match accurately enough the right tail of the wealth distribution that emerges from data. Several modelling attempts have been proposed: heterogeneous discounting and risk attitudes, the presence of large shocks in the income process, the presence of bequest motives, and health shocks (see De Nardi and Fella (2017)). However, the one I rely on in this work was proposed by Quadrini (2000), Cagetti and Nardi (2006), and Buera and Shin (2013) among others. They introduce entrepreneurs in the model whose presence generates higher wealth concentration. This feature is due essentially to two channels: on the one hand, the heterogeneity in returns due to credit constraints, and on the other hand, the need to save to overcome credit constraints itself. It is interesting to notice that entrepreneurs in these models are represented using Lucas' span-of-control framework. Moreover, the importance of entrepreneurs is not only theoretical but also empirical: despite being a small fraction of the population (about 8%) they hold a large share of wealth (about 40%) (see De Nardi and Fella (2017)). Analogously, Smith et al. (2019) highlight the fact that entrepreneurial skills are a crucial component of the firm's performance and that in terms of income, entrepreneurs are a significant fraction of the top earners. In this sense, studying how technological progress affects entrepreneurial returns is important to evaluate top income and wealth dynamics.

The dynamic model described in the second part of Section 3.2 builds on these types of models in order to have a characterisation of the top wealth concentration. Similar and recent works are those by Salgado (2019) and Koru (2020). The former studies how the decline in capital price increased firms' ability to hire high-skilled workers who would otherwise have become entrepreneurs, while the latter focuses on the impact of automation on top wealth inequality. This last one belongs to the literature studying the determinants of the increase in wealth inequality documented by Saez and Zucman (2016) and Kopczuk (2015).

The second strand of literature related to this chapter is that of skill-biased technological change, Violante (2008) is a succinct but enjoyable review of its main concepts. Traditionally (Solow (1957)), technological change was seen as factor-neutral scaling of the production technology, and, therefore,¹.

¹Another view, more natural if one adopts good old activity analysis Koopmans (1953), suggests that technical

A rise in the price of skilled labour, empirically registered since the Seventies Acemoglu and Autor (2011) Katz and Murphy (1992), suggests that technological change is factor-biased, and, in particular, skilled labour-biased: a factor-neutral technical change would not affect relative prices. Several formulations of skill-biased technological change exist, but I report the two that are more relevant for this work. The first one (Krusell et al. (2000)) is based on the observation that the price of capital has been decreasing throughout the years, leading to an increase in its use in production. Since capital is likely to be complementary with skilled labour, this triggered the increase in skilled labour demand and skill premium. Salgado (2019) uses this version of skill-biased technological change to explain the decline in the U.S. entrepreneurship rate: cheaper capital increased firms' ability to hire high-skilled individuals who would have otherwise attempted entrepreneurship. A second view (Nelson and Phelps (1966)) sees skilled agents as better suited in times of rapid technological change: they learn faster how to properly use new technologies, at least in the early stages of new technological waves. This view could be used to interpret the entrepreneurial technology that I adopt as a short-run version of a more general process of technological waves (Galor and Moav (2000)), where initially, the most skilled extract higher profits before slower agents learn how to use new technologies.

A final strand of research related to this paper is the one documenting the decline in entrepreneurship rate in the U.S. (Salgado (2019)) and the related decline in business dynamism (Akcigit and Ates (2021) Haltiwanger (2022)). Entrepreneurs are vital for the birth of new firms, job creation, and aggregate productivity. Salgado (2019) finds that, depending on the definition adopted, the share of entrepreneurs declined on average by 5% with respect to the value in 1985. Another relevant empirical fact related to the decline in business dynamism is the divergence in the productivity distribution of firms as documented in Andrews, Criscuolo, and Gal (2016) and reported in Figure 3.3. The modelling device I adopt also speaks towards these dimensions: Poschke (2018) already showed that his way of modelling the interaction between the general technological level and entrepreneurs' skills explains well the divergence in the distribution of firm size across different stages of development, and I aim at using it in the dynamic model as well. Moreover, in the simple static model of Section 3.2, an increase in technology reduces the mass of entrepreneurs in the economy: in general equilibrium, when firms become larger due to the increase of technology, also their labour demand increases and this, in turn, raises the equilibrium wage and, consequently, the outside option from entrepreneurship.

change is a move along a given isoquant, i.e., a change of the way of producing a given unit of output

3.2 Theoretical framework

In this section, I present the theoretical framework used in the analysis. In the first part, I describe a static version of the model, which can be considered as a building block of the dynamic one. I use the static version to show how exogenous technological progress can impact income inequality. This model's interaction between general technology and individual abilities generates skill-biased technological change. While technological progress improves every agent's potential firm productivity, the most skilled individuals extract more significant benefits. This mechanism translates into a higher concentration of income. I run a simple counterfactual (being aware that it is a theoretical exercise and not a comprehensive explanation of the phenomena) to show that technological progress, proxied by Total Factor Productivity growth, explains most of the increase in top income inequality. In the second section, I describe a dynamic version of the model that I still need to solve completely. This model aims to focus on the ability of skill-biased technological change to generate an increase in the concentration of wealth inequality.

3.2.1 A static model

The aim of this section is to provide a stylized description of the relationship between technological change and top income inequality. To do so, I introduce exogenous economic growth into a standard span of control model à la Lucas (1978). The main assumption is borrowed by Poschke (2018), who assumes that general technological changes affect agents differently based on their general ability: while technological progress improves everyone's productivity, the most skilled individuals extract more significant gains. This assumption and its micro-foundations are described in Section 3.A.

The environment

In this section, I describe a simple model of a static economy with no capital. The economy is populated by agents who must decide whether to become entrepreneurs and run their own firms or to be workers. This choice depends on their skill level z , which follows a distribution $F(z)$ with support $[z_{min}, +\infty)$ ².

If an agent decides to become a worker, she supplies one unit of labour for a wage w , while if she opts for entrepreneurship, she runs a firm with technology:

$$\max_l M^z l^\gamma, \quad (3.1)$$

where the firm productivity is given by two components: the skill level of the agent, z , and the general technology level available in the economy M . The objective of entrepreneurs is to maximize profits by solving:

$$\max_l M^z l^\gamma - wl, \quad (3.2)$$

The solution of this problem is trivial and leads to a labour demand and a profit function that depends on M and z themselves:

$$l^d = \left(\frac{\gamma M^z}{w} \right)^{\frac{1}{1-\gamma}} \quad \text{and} \quad \pi(M, z) = (1 - \gamma) M^{\frac{z}{1-\gamma}} \left(\frac{\gamma}{w} \right)^{\frac{\gamma}{1-\gamma}} \quad (3.3)$$

²Later, the distribution is assumed to be exponential with rate λ , with $z_{min} = 0$.

The occupational choice is described by the following binary variable ($e = 1$ if entrepreneur and $e = 0$ if worker):

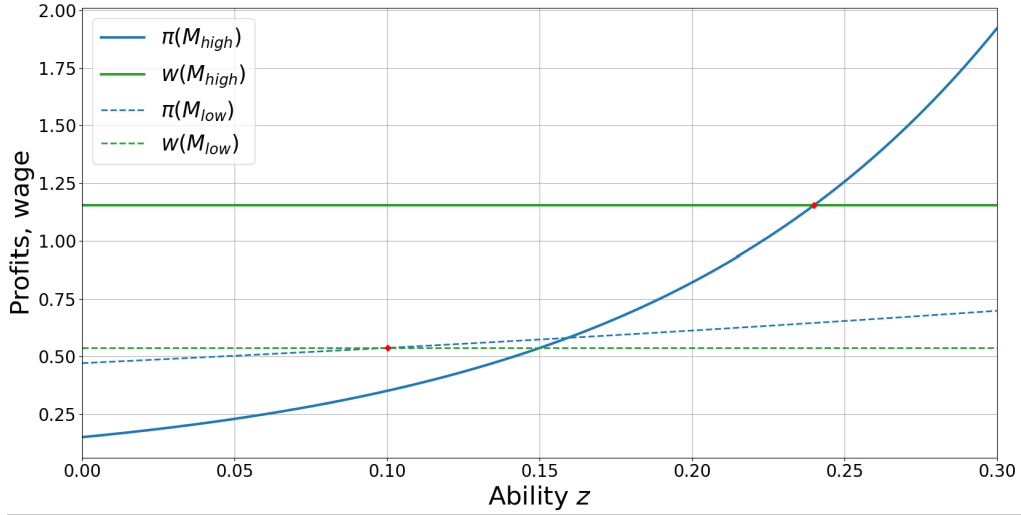
$$e(z) = \begin{cases} 1 & \text{if } \pi(M, z) > w \\ 0 & \text{if } \pi(M, z) \leq w \end{cases} \quad (3.4)$$

The equilibrium

An equilibrium in this economy is a couple (z^*, w^*) , where the former is the ability threshold below which agents become workers, while the latter is the equilibrium wage³. These quantities can be found by solving the two following conditions:

1. $\pi(M, z) = w$ (occupational choice)
2. $\int_{z_{min}}^{z^*} f(z) dz = \int_{z^*}^{+\infty} l^d(w; M, z) f(z) dz$ (labour market clearing)

Figure 3.4: Characterisation of equilibrium at two different technology level, M_{low} and M_{high}



From condition 1. it is possible to express the value of the cutoff z^* as a function of the equilibrium wage:

$$w = (1 - \gamma)M^{\frac{z}{1-\gamma}} \left(\frac{\gamma}{w}\right)^{\frac{\gamma}{1-\gamma}} \Rightarrow z^*(w) = \frac{1}{\log(M)} \left(\log(w) - \gamma \log(\gamma) - (1 - \gamma) \log(1 - \gamma) \right) \quad (3.5)$$

By considering a specific distribution for the agents' skills, it is possible to get an analytical version of condition 2. If we assume that skills are distributed as an exponential with rate λ , then we obtain:

$$1 - e^{-\lambda z^*} = \int_{z^*}^{+\infty} \left(\frac{\gamma M^z}{w}\right)^{\frac{1}{1-\gamma}} \lambda e^{-\lambda z} dz, \quad (3.6)$$

³The economy is fully characterised by three parameters: the span of control in entrepreneurial technology, γ , the technological level, M , and the parameter of the skill distribution, λ .

which, thanks to equation 3.5, can be solved for w (eventually for various levels of M). In particular, the right-hand-side of Equation 3.6 is equal to:

$$\left(\frac{\gamma}{w}\right)^{\frac{1}{1-\gamma}} \lambda \int_{z^*}^{+\infty} M^{\frac{z}{1-\gamma}} e^{-\lambda z} dz = \left(\frac{\gamma}{w}\right)^{\frac{1}{1-\gamma}} \lambda \int_{z^*}^{+\infty} e^{z \frac{\log(M)}{1-\gamma}} e^{-\lambda z} dz. \quad (3.7)$$

This integral converges as long as $\log(M) - \lambda(1 - \gamma) < 0$, i.e., if $M < \exp(\lambda(1 - \gamma))$. In this case, Equation 3.6 becomes:

$$-(1 - e^{-\lambda z^*(w)}) + \lambda \left(\frac{\gamma}{w}\right)^{\frac{1}{1-\gamma}} \frac{1 - \gamma}{\log(M) - \lambda(1 - \gamma)} \exp\left[\frac{\log(M) - \lambda(1 - \gamma)}{1 - \gamma} z^*(w)\right] = 0, \quad (3.8)$$

which is the Excess Demand for Labour at wage w . Figure 3.5 represents Labour Excess Demand for different values of the general technological level M .

Figure 3.5: Labour excess demand for different technological levels





Figure 3.6: Equilibrium wage and occupational choice cutoff at different levels of technology M 

Table 3.1: Parameters of the static model

Parameter	Value
exponential rate (λ)	10
span of control (γ)	0.6
technology (M)	(1, 30]

A characterisation of top income inequality

Assuming that z is a random variable, we can also compute the distribution of profits. It should be noted that this is only an approximation of the overall income distribution in the model, as workers' earnings (wages) are excluded from the computations. Since the main focus is on the share of income accruing to the top earners, who by construction are entrepreneurs, focusing on the distribution of profits is a reasonable approximation. For convenience, let us define:

$$c := (1 - \gamma) \left(\frac{\gamma}{w}\right)^{\frac{\gamma}{1-\gamma}} \quad \text{and} \quad \beta := \frac{1}{1 - \gamma} \quad \Rightarrow \quad \pi(M, z) = cM^{\beta z}. \quad (3.9)$$

If we assume that skills are distributed exponentially with a rate of λ , then the distribution of profits can be obtained as follows:

$$\begin{aligned} \mathbb{P}(\pi \leq \hat{\pi}) &= \mathbb{P}(cM^{\beta z} \leq \hat{\pi}) = \mathbb{P}(\beta z \log(M) \leq \log\left(\frac{\hat{\pi}}{c}\right)) = \\ &= \mathbb{P}\left(z \leq \log\left(\frac{\hat{\pi}}{c}\right) \frac{1}{\beta \log(M)}\right) = \int_0^{\hat{z}} \lambda e^{-\lambda z} dz = e^{-\lambda \hat{z}} \Big|_0^{\hat{z}} = 1 - e^{-\lambda \hat{z}} = 1 - \left(\frac{\hat{\pi}}{c}\right)^{\frac{-\lambda}{\beta \log(M)}}. \end{aligned} \quad (3.10)$$

Following a conventional notation, I call

$$\frac{-\lambda}{\beta \log(M)} = -\alpha + 1. \quad (3.11)$$

Then, it can be shown that profits follow a Pareto distribution with shape parameter α and scale parameter c . If we change variable, $\hat{\pi} = x$, and take the derivative, we obtain the density function:

$$f(x) = \frac{\alpha - 1}{c} \left(\frac{x}{c}\right)^{-\alpha}. \quad (3.12)$$

With this notation, calculations are easier, and we can find the q -th quantile of the distribution, i.e.,

$$x_q \text{ s.t. } \mathbb{P}(\pi \geq x_q) = q, \quad \text{by solving: } \int_{x_q}^{+\infty} \frac{\alpha - 1}{c} \left(\frac{x}{c}\right)^{-\alpha} dx = q. \quad (3.13)$$

The q -th quantile is, therefore,

$$\left(\frac{x_q}{c}\right)^{-\alpha+1} = q \quad \Rightarrow \quad x_q = cq^{\frac{1}{1-\alpha}}. \quad (3.14)$$

Now it is possible to compute the fraction of entrepreneurial income accruing to the agents belonging to the q -th quantile⁴. Let us define such a share as follows:

$$S(q) = \frac{\int_{x_q}^{+\infty} x \frac{\alpha-1}{c} \left(\frac{x}{c}\right)^{-\alpha} dx}{\mathbb{E}(\pi(x))} \quad (3.15)$$

This is a standard result for Pareto distributions, and relatively easy calculations show that:


$$S(q) = q^{\frac{\alpha-2}{\alpha-1}}, \quad (3.16)$$

which implies that if α increases, $S(q)$ decreases. Substituting the original coefficients of the profit function 3.9, we obtain a relationship between exogenous technological change, M , and top-income inequality $S(q)$, namely:

$$S(q) = q^{\frac{\lambda - \beta \log(M)}{\lambda}}. \quad (3.17)$$

The relationship between technological level and top income inequality is plotted in Figure 3.7, which shows a positive association. The pattern is due to the fact that when technology increases, it benefits those who are more skilled to a greater extent. From an economic viewpoint, technological progress increases entrepreneurial returns at an increasing rate with their skills.

⁴Please notice that I make a slight simplification here: first of all, this is not the whole income distribution since I am ignoring the wage-earners; secondly I am ignoring the possibility that there might be a mass of profit earners between the occupational cut-off z^* and the scale parameter c . These assumptions are innocuous since I am focusing only on the right tail of the profit distribution, which by construction is made by profit-earners only.


Figure 3.7: Top income shares at different levels of technology M 

Finally, Figure 3.8 shows a simple counterfactual exercise: I assume that the initial share of income earned by, respectively, the top 1% and the top 10% earners in the distribution are those of the United States in 1980 (respectively, 10.4% and 33.9%). Then, I assume that technology M grows at the same average pace as the United States Total Factor Productivity over the same period (31.6%) and plot the evolution of the income shares⁵. Figure 3.8 can be compared with Figure 3.1: while both the simulated top 10% and top 1% income shares in the final year (i.e., when the whole growth has been applied) underestimate the actual values, the patterns are qualitatively similar to the actual values⁶.

⁵Data on income inequality come from the World Inequality Database, data on TFP from Penn World Table

⁶Actual top 1% in 2020 earns 17.9% of total income, against the 14% of the simulated data; Actual top 10% in 2020 earns 44.4% of total income, against the 40% of the simulated data).

Figure 3.8: Simple counterfactual exercise



Evolution of top income shares assuming that technology M grows at the US TFP rate over 1980-2020 (33.6%)

3.2.2 A dynamic model

Overview

The aim of the model is to study how the introduction of skill-biased technological change in entrepreneurial technology (as defined in Poschke (2018)) allows us to explain the stylised facts described in the previous section: the increase in wealth concentration, the widening gaps in firms productivity and size distributions, and the decline in entrepreneurship.

I nest the entrepreneurial technology assumption in an Aiyagari economy (Aiyagari (1994)) with entrepreneurs (Quadrini (2000) and Cagetti and Nardi (2006)). The presence of entrepreneurs is crucial for understanding wealth accumulation, while the assumptions on technology are needed to show the increase in concentration. In what follows, I borrow the notation and part of the model from Buera and Shin (2013) and Poschke (2018).

Demographics and Preferences

The economy is populated of infinitely-lived individuals that are heterogeneous with respect to their wealth a (I denote the assets space by \mathcal{A}) and their intrinsic ability $z \in Z$. While a depends on the optimal saving choices of the agents, z comes from an invariant distribution with support Z . In every period, an agent is endowed with a draw from this distribution; in the next period, he has a probability ψ to maintain the previous period's ability and with a probability $1 - \psi$ to draw (independently from the current level) another value of z . The size of the population is normalised to one, and we denote by $\mu(z)$ the fraction of the population with ability z . I use $\Gamma_t(a, z)$ to denote the joint distribution of the population with wealth a and ability z at time t .

Individual preferences for streams of consumption at time $t = 0$ are given by the standard expected

utility:

$$U = \mathbb{E}_0 \sum_{t=0}^{+\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma}$$

Technology and Markets

In every period, the agents decide whether to start a firm with or to work for a wage w . In the former case, the individual becomes an entrepreneur and produces an output y by hiring labour n and renting capital k . The firm's technology will depend on the individual's ability z :

$$y = M^z (k^\alpha n^{1-\alpha})^\gamma. \quad (3.18)$$

where M is a measure of the general technology level, and γ is the span-of-control parameter (see Lucas (1978) for early adoption of this modelling technique). If he decides to be a worker, he will obtain labour earnings that are proportional to his skills, wz . The labour market is assumed to be perfectly competitive. Capital k depreciates at a rate δ and is rented at a rate r from a competitive intermediary, which pools agents' savings a . Agents are liquidity constrained and therefore cannot borrow, i.e. $a \geq 0$.

Another typical assumption in these types of models, which traces back to Evans and Jovanovic (1989), is that entrepreneurs can rent capital up to a multiple of their own wealth: $k \leq \lambda a$. If λ approaches to 1, the firm's capital is entirely financed by the entrepreneur, and relaxing λ amounts to relaxing the degree of credit frictions. This assumption is a simplified version of models where credit constraints are determined endogenously and that rely on the assumption of imperfect contract enforceability (see Cagetti and Nardi (2006)) for an important application of this concept to the entrepreneurs' problem). In such models, the creditors cannot force debt repayments, and the debtor (the entrepreneur in our case) will repay the borrowed capital only if he finds it convenient.

This assumption is crucial to obtain returns heterogeneity: without collateral constraints, every entrepreneur would invest up to the level that equalises the return with the economy interest rate r .

The other crucial assumptions are borrowed by Poschke (2018), and they refer to the entrepreneurial technology and the workers' earnings. The classical assumption in occupational choice models, used, for example, to study firm size distribution (Lucas (1978)), or more generally firm dynamics Hopenhayn (1992), is that a firm's productivity is linear in the owner's ability. In my model, the firm productivity is made of two components: M and z . The former is common to all firms and can be considered a measure of the general technology available in the economy. We can represent technological progress as an (exogenous) increase in M . The latter is a measure of the entrepreneurial ability. The crucial point about this formulation of the firm's productivity is that an increase in M will favour the more skilled entrepreneurs. In appendix 3.3, I describe the properties of this assumption in greater detail.

Timing and Agents' Problem

At time $t = 0$ the problem of the agents is the following:

$$\max_{\{c_t, a_{t+1}\}_t} \mathbb{E}_0 \sum_{t=0}^{+\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma} \quad (3.19)$$

$$\begin{aligned} \text{s.t. } c_t + a_{t+1} &= \max_{e_t \in \{0,1\}} \{e_t \pi_t + (1 - e_t) w_t z_t\} + (1 + r_t) a_t \\ a_t &\geq 0 \quad \forall t, \end{aligned}$$

where e_t represents the occupational choice ($e = 1$ if entrepreneur and $e = 0$ if worker) and π_t is the period t profit function of the entrepreneur. Profits come from a standard profit maximisation problem where the entrepreneur hires labor n at wage w and rents capital k at a rate $R := r + \delta$ given the collateral constraint:

$$\max_{n, k \leq \lambda a} \{(M^z)^{1-\gamma} (k^\alpha n^\beta)^\gamma - Rk - wn\} \quad (3.20)$$

The value function of this problem will be contingent on the state variables (a, z) , and in particular on the entrepreneur being collateral constrained or not (the derivations can be found in the Appendix 3.B):

$$\pi(a, z; w, R) = \begin{cases} \pi^u & \text{if } M^z \leq h(a; \lambda) \\ \pi^c & \text{if } M^z > h(a; \lambda), \end{cases}$$

where $h(a; \lambda)$ is a function that is increasing both in the available wealth a and in the collateral constraint parameter λ . After finding profits, we can calculate the occupational choice of the individual (again, completely described by the state variables (a, z)) by comparing profits, which can be constrained or not, with labour earnings:

$$e(a, z; w, R) = \begin{cases} 1 & \text{if } \pi(a, z; w, R) > wz \\ 0 & \text{if } \pi(a, z; w, R) \leq wz \end{cases}$$

We can represent the problem recursively (and we need to, in order to solve it): at the beginning of each period, after observing the state (a, z) , each individual decides whether to be a worker and get a labour income wz , or to be an entrepreneur and start a firm with the technology (3.18).

$$V(a, z) = \max_{c, a', e} \{u(c) + \beta \mathbb{E}_z z V(a', z')\} \quad (3.21)$$

$$\text{s.t. } c + a' = \max_e \{\pi, wz\} + (1 + r)a \quad (3.22)$$

$$a \geq 0 \quad (3.23)$$

Equilibrium Definition

Given an initial distribution of wealth and abilities, $\Gamma_0(a, z)$, a competitive equilibrium for this economy consists of sequences of decisions $\{c_t(a_t, z_t), a_{t+1}(a_t, z_t), l_t(a_t, z_t), k_t(a_t, z_t)\}_{t=0}^{+\infty}$, prices $\{w_t, r_t\}_{t=0}^{+\infty}$, and distributions $\{\Gamma_t(a_t, z_t)\}_{t=1}^{+\infty}$ such that :

1. Given prices $\{w_t, r_t\}_{t=0}^{+\infty}$ and the state variables (a_t, z_t) , $\{c_j(a_t, z_t), a_{j+1}(a_t, z_t), l_j(a_t, z_t), k_j(a_t, z_t)\}_{j=t}^{+\infty}$ solve the individuals' problems 3.19 for all $t \geq 0$

2. Markets clear⁷. In particular:

$$\sum_{z \in Z} \mu(z) \left\{ \int_E l(a, z; w_t, r_t) d\Gamma_t(a|z) - \int_{E^c} z d\Gamma_t(a|z) \right\} = 0 \quad (\text{Labour Market})$$

$$\sum_{z \in Z} \mu(z) \left\{ \int_E k(a, z; w_t, r_t) d\Gamma_t(a|z) - \int_{E \cup E^c} a d\Gamma_t(a|z) \right\} = 0 \quad (\text{Capital Market})$$

3. the joint distribution evolves according to the mapping:

$$\Gamma_{t+1}(a|z) = \psi \int d\Gamma_t(v|z) du + (1 - \psi) \sum_{\hat{z} \in Z} \mu(\hat{z}) \int d\Gamma_t(v|\hat{z}) du$$

Preliminary computations

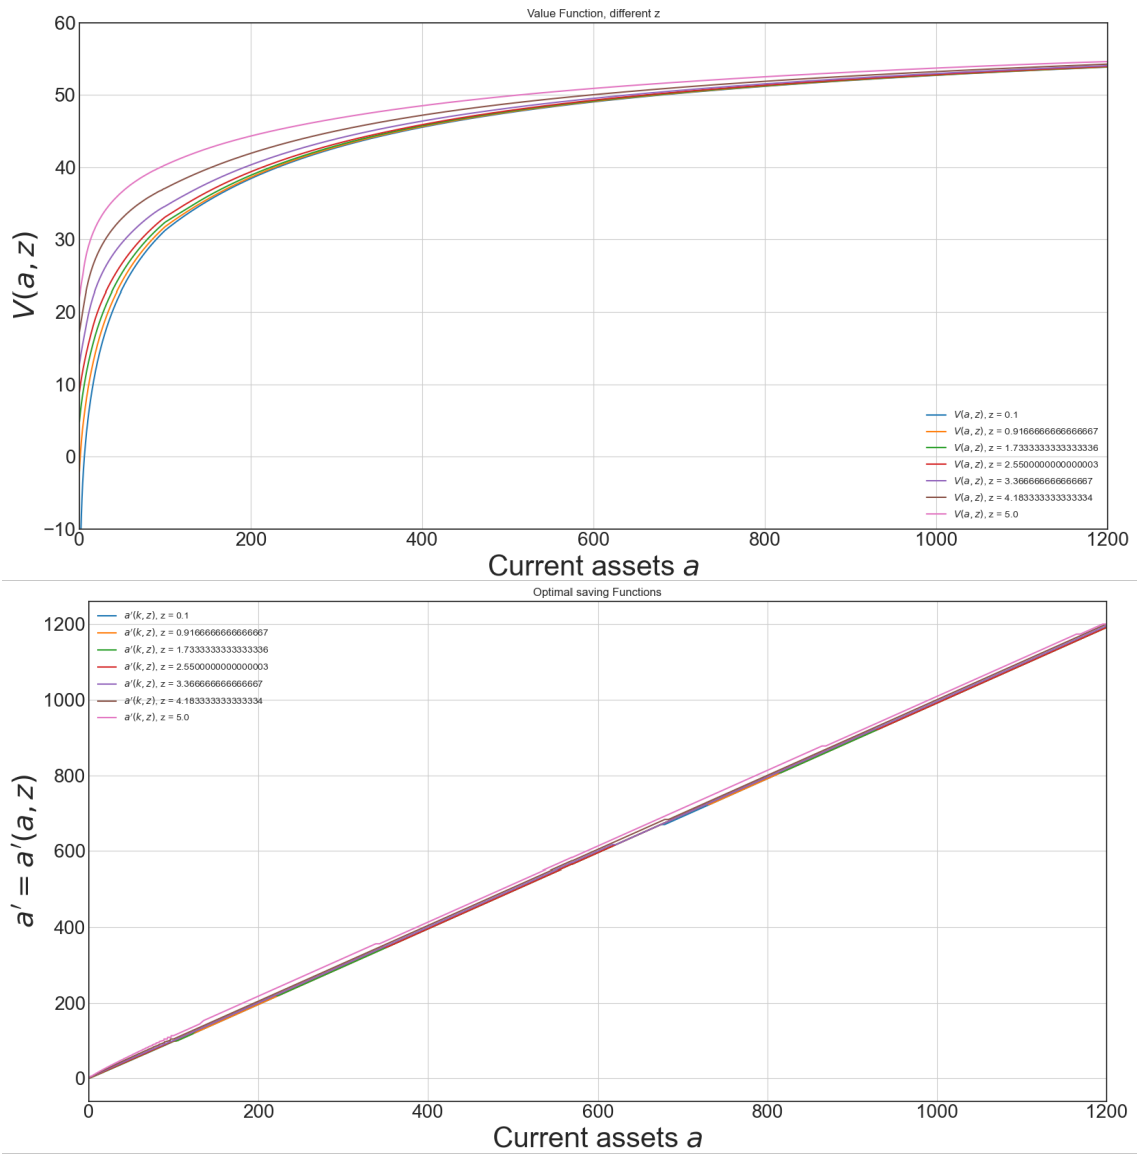
In Figure 3.9, I report the value function and the policy function for optimal savings of the agents in the dynamic model. This are obtained by value function iteration on a grid of state variables $\mathcal{A} \times Z = [0, 1200] \times [0.1, 5]$.

Table 3.2: Parameters of the dynamic model

Parameter	Value
discount factor (β)	0.97
risk-aversion (σ)	1.5
span of control (γ)	0.8
capital share (α)	0.33
collateral constraint (λ)	1.5
capital depreciation (δ)	0.06
prob. of retain ability (ψ)	0.894

⁷Notation: occupational choice partitions the state space $\mathcal{A} \times Z$ into two regions E and E^c which represent, respectively, the agents who become entrepreneurs and those who become workers

Figure 3.9: Preliminary computations



Value Function $V(a, z)$ of the agent (above) and (savings) policy function $a' = a'(a, z)$ (below) over the asset space for different skills levels.

3.3 Conclusions

In this short note, I briefly touch on a few of the several empirical patterns studied in the macroeconomic literature in the last decades: the increase in top income and top-wealth inequality (see Kopczuk (2015) and Saez and Zucman (2016)), the decline in the entrepreneurship rate (Haltiwanger (2022), Kozeniauskas (2018), and Salgado (2019)), and the so-called divergence in the firm distribution (Andrews, Criscuolo, and Gal (2016), Akcigit and Ates (2021)). Please notice that this is a simple exposition of a possible modelling device rather than a full explanation of these phenomena. In the first part of this chapter, I build a simple model to show the relationship between skill-biased technological change and top-income inequality and conduct a counterfactual experiment that qualitatively replicates the increase in top-income concentration by assuming that technological change can be proxied by Total Factor Productivity. When technological progress occurs, the entrepreneurs closer to the frontier become increasingly more productive, and income concentration, which for entrepreneurs is given by their profits, increases. From an economic viewpoint, technological progress increases entrepreneurial returns at an increasing return with their skills. A by-product of the model is that, in general equilibrium, when firms become larger due to the increase of technology, their size increases and, therefore, also their labour demand. This raises the equilibrium wage and, consequently, the outside option from entrepreneurship. Therefore, the mass of entrepreneurs in the economy decreases, in line with the findings by Salgado (2019). The technology specification also implies a Pareto distribution of entrepreneurial incomes, without assuming a Pareto distribution in their skills, as it was done in the original article by Lucas (1978). In the second part, I describe a more elaborate model that I plan to use to study the effect of technological change on wealth inequality, but that I still need to fully solve. The dynamic model, in principle, rests on the same mechanisms on which the static model is based, but it uses entrepreneurs to generate a level of wealth concentration that is consistent with the data. The idea, if the model can be solved, is to study the impact of technological change on wealth concentration at the top when mediated by entrepreneurial skill-biased technological change.

In this appendix, I report a number of calculations to support some results claimed in the main text. In the first section, I solve the constrained profit maximisation problem faced by the entrepreneur, while in the second, I give some details about the occupational choice decision.

3.A Entrepreneurial Technology

In this section I briefly describe the assumption at the base of the entrepreneurial technology and we can think of them as a sort of microfoundations for it. This paragraph relies heavily on Poschke (2018).

$$y = X^\gamma \quad \text{where} \quad X = \left(\int_0^H x_j^{\frac{\rho-1}{\rho}} dj \right)^{\frac{\rho}{\rho-1}} \quad \text{and} \quad x_j = n_j^{(1-\alpha)} k_j^\alpha \quad (3.24)$$

Where y is the final output, X represents the intermediates, $\gamma < 1$ is the span of control parameter. The variable x_j is the intermediate tasks employing capital k and labour n ; ρ is the elasticity of substitution between tasks and, finally, H is the firm's technological level. H can be thought as the number of tasks supervised, a measure of complexity of the firm's activity, or its productivity. As previously mentioned, H is made of two components: the general level of technology M and the agent's ability z : $H = H(M, z)$. The following assumptions are imposed to this function, and I will briefly describe them.

- i) $\partial H(z, M)/\partial z > 0$
- ii) $\partial H(z, M)/\partial M > 0$
- iii) the elasticity of $H(z, M)$ w.r.t. M does not depend on M
- iv) the elasticity of $H(z, M)$ w.r.t. \bar{M} increases in z
- v) the elasticity of $H(z, M)$ w.r.t. M is convex in z

Assumption i) and ii) are straightforward and, respectively, tell that firms' productivity increase when general technology increase and that more skilled entrepreneurs can manage bigger firms. Assumption iii) is a technical one that helps with tractability. Assumption iv) is the one governing skill biased technological change: it says that more able entrepreneurs benefit relatively more from a general technological advancement. Finally, assumption v) is crucial for generating the occupational choices presented in the model, and rely on profits being convex in H and z . The simplest function fulfilling all these property is the exponential M^z adopted in the main text.

3.B Constrained Profit Maximization

If the individual becomes an entrepreneur maximises the within-period profits by hiring labour n and renting capital k at a rate $R := r + \delta$:

$$\max_{n, k, \nu} \{M^z (k^\alpha n^{(1-\alpha)})^\gamma - wn - Rk\} \quad (3.25)$$

Profits will depend exclusively on prices and the state variables of the problem, so that the profit function will be $\pi(a, z; w, r)$. The problem is standard, we just need to take into account the collateral constraint $k \leq \lambda a = \hat{k}$. To simplify the algebra let's rename the parameters: $H := M^z$, $\mu := \alpha\gamma$, $\theta := (1 - \alpha)\gamma$, $D := 1 - \mu - \theta = 1 - \gamma$. The Lagrangean of the problem becomes:

$$\mathcal{L}(k, n, \nu) := Hk^\mu n^\theta - wn - Rk - \nu(k - \hat{k}) \quad (3.26)$$

The unconstrained case: the optimal capital level is below the collateral constraint, $k^* < \hat{k}$, and the multiplier $\nu = 0$ and the profit maximisation is the standard one. Since the derivation is standard I only report the unconditional factor demands:

$$n^*(z; w, r) = H^{\frac{1}{D}} \left(\frac{\theta}{w} \right)^{\frac{1-\mu}{D}} \left(\frac{\mu}{R} \right)^{\frac{1-\mu}{D}}, \quad k^*(z; w, r) = H^{\frac{1}{D}} \left(\frac{\theta}{w} \right)^{\frac{\theta}{D}} \left(\frac{\mu}{R} \right)^{\frac{1-\theta}{D}} \quad (3.27)$$

and the profit function:

$$\pi^u := \pi(z; w, r) = (1 - \mu - \theta) H^{\frac{1}{D}} \left(\frac{\theta}{w} \right)^{\frac{\theta}{D}} \left(\frac{\mu}{R} \right)^{\frac{\mu}{D}} \quad (3.28)$$

The constrained case: $\nu^* > 0$ and $k^* \geq \hat{k}$, the entrepreneur is constrained and there is a wedge between the rental rate R and the marginal productivity of capital. The optimality conditions are:

$$[n] : \theta H k^\mu n^{\theta-1} = w, \quad [k] : \mu H k^{\mu-1} n^\theta - R - \nu \leq 0 \quad (3.29)$$

Therefore the unconditional demands depend also on a and are:

$$n^*(a, z; w, r) = \left(\frac{\theta H}{w} \right)^{\frac{1}{1-\theta}} (\lambda a)^{\frac{\mu}{1-\theta}}, \quad k^*(a, z; w, r) = H^{\frac{1}{D}} \left(\frac{\theta}{w} \right)^{\frac{\theta}{D}} \left(\frac{\mu}{R} \right)^{\frac{1-\theta}{D}} \quad (3.30)$$

and the profit function is:

$$\pi^c := \pi(a, z; w, r) = (1 - \theta) H^{\frac{1}{1-\theta}} (\lambda a)^{\frac{\mu}{1-\theta}} \left(\frac{\theta}{w} \right)^{\frac{\theta}{1-\theta}} - R(\lambda a). \quad (3.31)$$

Finally, from the multiplier ν^* being positive we can find a condition on H , that is M^z , which tells us when the entrepreneur is constrained given his wealth level:

$$[k] : \mu H (\lambda a)^{\mu-1} n^\theta - R = \nu^* > 0 \quad \iff \quad H > \left(\frac{R}{\mu} \right)^{1-\theta} \left(\frac{w}{\theta} \right)^\theta (\lambda a)^{1-\mu-\theta} =: H^* \quad (3.32)$$

This condition tells us that for given level of wealth a the most skilled will be constrained. In short we found that the profit function of an agent becoming an entrepreneur is:

$$\pi(a, z; w, R) = \begin{cases} \pi^u & \text{if } M^z \leq H^* \\ \pi^c & \text{if } M^z > H^* \end{cases}$$

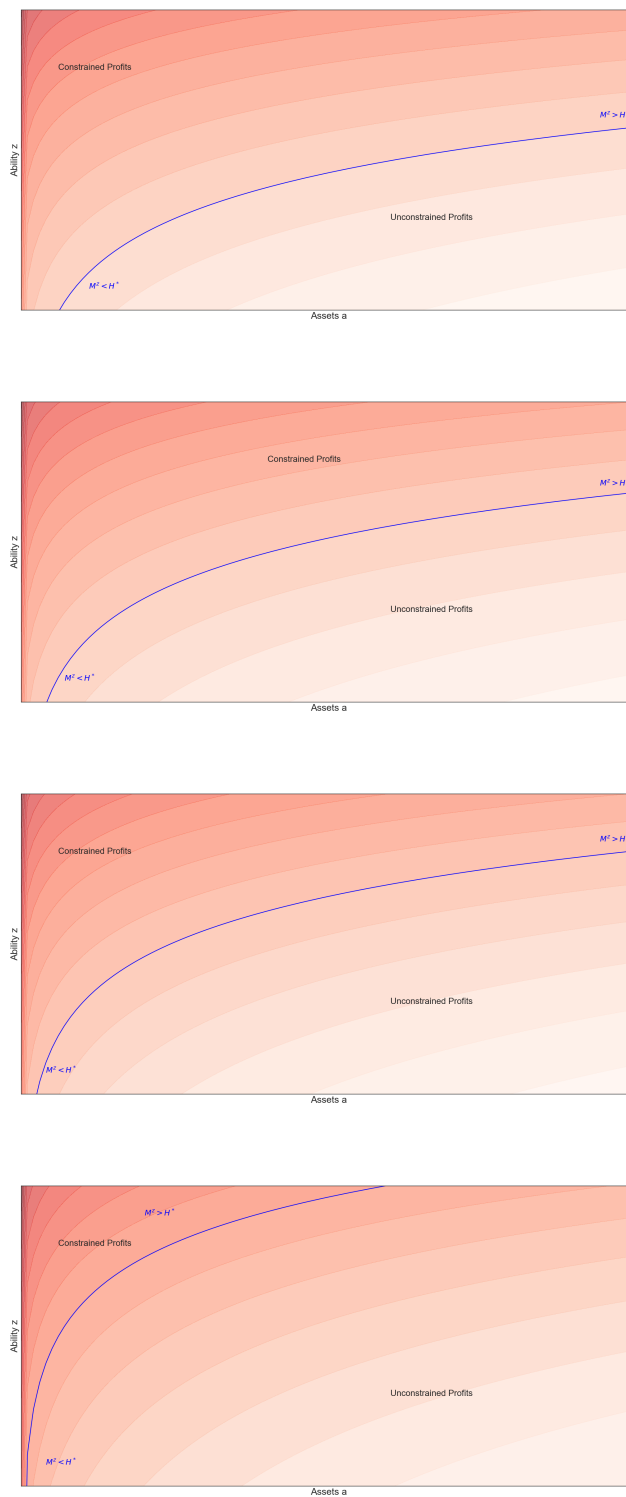


Figure 3.10: Top to bottom: decreasing credit frictions (increasing λ)

3.C Occupational Choice

In every period the agent compares potential profits $\pi(a, z; w, R)$, obtained in the previous section, with labour income wz . This comparison determines the binary variable $e(a, z; w, R)$, which represents the individual's occupational choice.

If the individual is not constrained, and this happens whenever

$$M^z \leq \left(\frac{R}{\mu}\right)^{1-\theta} \left(\frac{w}{\theta}\right)^\theta (\lambda a)^{1-\mu-\theta},$$

then he compares:

$$\pi^u \geq wz \iff (1 - \mu - \theta)H^{\frac{1}{b}} \left(\frac{\theta}{w}\right)^{\frac{\theta}{b}} \left(\frac{\mu}{R}\right)^{\frac{\mu}{b}} \geq wz$$

If instead he is constrained, i.e. whenever

$$M^z > \left(\frac{R}{\mu}\right)^{1-\theta} \left(\frac{w}{\theta}\right)^\theta (\lambda a)^{1-\mu-\theta},$$

then the comparison will be the following:

$$\pi^c \geq wz \iff (1 - \theta)H^{\frac{1}{1-\theta}} (\lambda a)^{\frac{\mu}{1-\theta}} \left(\frac{\theta}{w}\right)^{\frac{\theta}{1-\theta}} - R(\lambda a) \geq wz$$

In either case the positiveness of the parameters and the fact that profits are convex in the ability while earnings are linear, this equation will have two solutions, that we can denote as z_L and z_H . In particular, the individuals below and above such threshold will become entrepreneurs, while individuals in the middle will decide to become workers (see Figure 3.11). The more general relationship between credit frictions and occupational choice is represented in Figure 3.12: for low levels of wealth agents with high skill find convenient to remain workers, but as soon a certain level of wealth is accumulated they prefer to become entrepreneurs. We can express the occupational decision as a binary choice as follows:

$$e(a, z; w, R) = \begin{cases} 1 & \text{if } \pi(a, z; w, R) > wz \\ 0 & \text{if } \pi(a, z; w, R) \leq wz \end{cases}$$

In this way sequential budget constraint

$$\text{s.t. } c + a' = \max_e \{\pi, wz\} + (1 + r)a$$

can be rewritten as:

$$c(a, z) + a'(a, z) = e\pi(a, z; w, R) + (1 - e)wz + (1 + r)a.$$




Figure 3.11: Occupational Choice, given wealth.

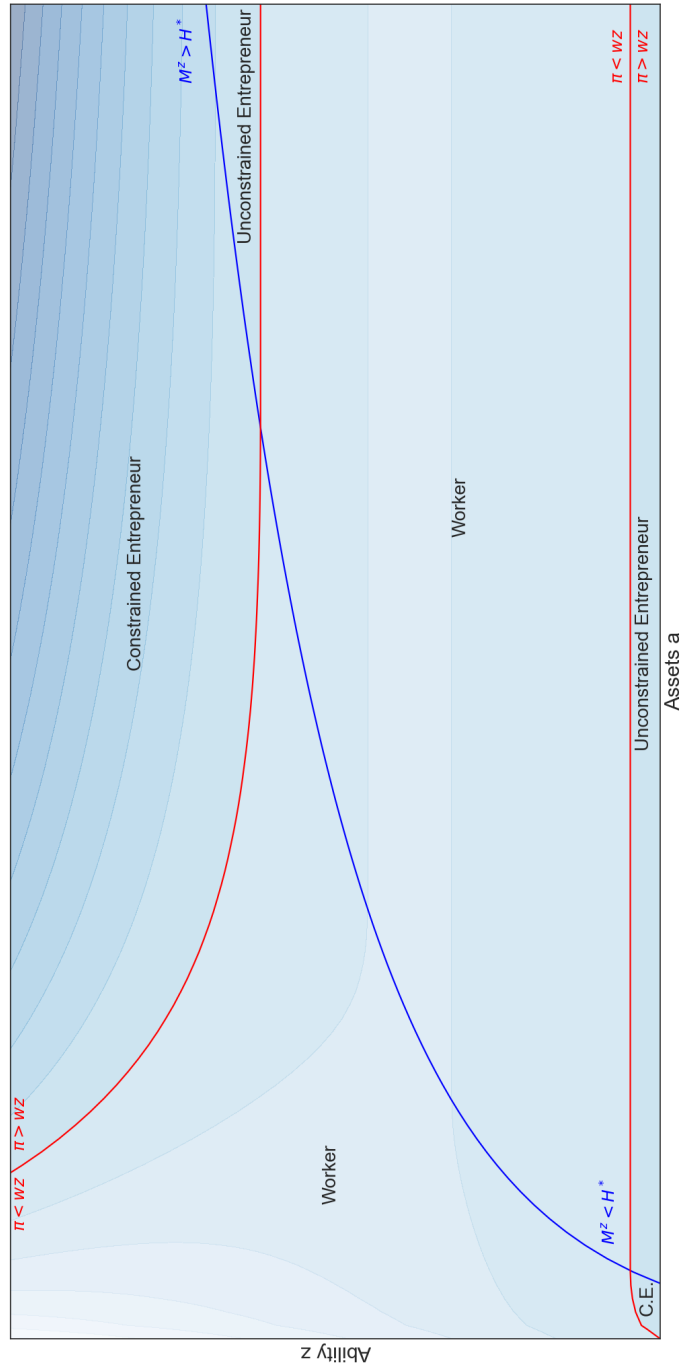


Figure 3.12: Occupational Choice over the whole state space (a, z) .




Figure 3.13: Occupational Choice at increasing levels of the wage w .

References

- Acemoglu, Daron (2002). “Technical Change, Inequality, and the Labor Market”. In: *Journal of Economic Literature* 40.1, pp. 7–72.
- Acemoglu, Daron and David Autor (2011). “Skills, Tasks and Technologies: Implications for Employment and Earnings”. In: vol. 4. *Handbook of Labor Economics*. Elsevier, pp. 1043–1171.
- Aiyagari, S. Rao (1994). “Uninsured Idiosyncratic Risk and Aggregate Saving”. In: *The Quarterly Journal of Economics* 109.3, pp. 659–684.
- Akcigit, Ufuk and Sina T. Ates (2021). “Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory”. In: *American Economic Journal: Macroeconomics* 13 (1).
- Andrews, Dan, Chiara Criscuolo, and Peter N. Gal (2016). “The Best versus the Rest”. In: 5.
- Bewley, Truman (1986). “Stationary monetary equilibrium with a continuum of independently fluctuating consumers”. In: *Contributions to mathematical economics in honor of Gérard Debreu* 79.
- Buera, Francisco J. and Yongseok Shin (2013). “Financial Frictions and the Persistence of History: A Quantitative Exploration”. In: *Journal of Political Economy* 121.2, pp. 221–272.
- Cagetti, Marco and Mariacristina De Nardi (2006). “Entrepreneurship, Frictions, and Wealth”. In: *Journal of Political Economy* 114.5, pp. 835–870.
- De Nardi, Mariacristina and Giulio Fella (Oct. 2017). “Saving and wealth inequality”. In: *Review of Economic Dynamics* 26, pp. 280–300.
- Evans, David S. and Boyan Jovanovic (1989). “An Estimated Model of Entrepreneurial Choice under Liquidity Constraints”. In: *Journal of Political Economy* 97.4, pp. 808–827.
- Gabaix, Xavier (Feb. 2016). “Power Laws in Economics: An Introduction”. In: *Journal of Economic Perspectives* 30 (1), pp. 185–206.
- Galor, Oded and Omer Moav (2000). “Ability-Biased Technological Transition, Wage Inequality, and Economic Growth”. In: *The Quarterly Journal of Economics* 115.2, pp. 469–497.
- Haltiwanger, John (2022). “Entrepreneurship in the twenty-first century”. In: *Small Business Economics*, pp. 1–14.
- Hopenhayn, Hugo A. (1992). “Entry, Exit, and firm Dynamics in Long Run Equilibrium”. In: *Econometrica* 60.5, pp. 1127–1150.
- Hubmer, Joachim, Per Krusell, and Anthony A. Smith. (May 2021). “Sources of US Wealth Inequality: Past, Present, and Future”. In: *NBER Macroeconomics Annual* 35, pp. 391–455.
- Huggett, Mark (1993). “The risk-free rate in heterogeneous-agent incomplete-insurance economies”. In: *Journal of economic Dynamics and Control* 17.5–6, pp. 953–969.
- Imrohoroğlu, Ayşe (1989). “Cost of business cycles with indivisibilities and liquidity constraints”. In: *Journal of Political economy* 97.6, pp. 1364–1383.
- Katz, L. F. and K. M. Murphy (Feb. 1992). “Changes in Relative Wages, 1963-1987: Supply and Demand Factors”. In: *The Quarterly Journal of Economics* 107 (1), pp. 35–78.
- Kaymak, Barış and Markus Poschke (Feb. 2016). “The evolution of wealth inequality over half a century: The role of taxes, transfers and technology”. In: *Journal of Monetary Economics* 77, pp. 1–25.
- Koopmans, Tjalling C (1953). “Activity analysis and its applications”. In: *The American Economic Review* 43.2, pp. 406–414.
- Kopczuk, Wojciech (Feb. 2015). “What Do We Know about the Evolution of Top Wealth Shares in the United States?” In: *Journal of Economic Perspectives* 29 (1), pp. 47–66.
- Koru, Ömer Faruk (Nov. 2020). “Automation and Top Wealth Inequality”. In.
- Kozeniauskas, Nicholas (2018). “What’s Driving the Decline in Entrepreneurship?” In: *Unpublished paper. New York University, New York, NY*.

- Krusell, Per et al. (2000). “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis”. In: *Econometrica* 68.5, pp. 1029–1053.
- Lucas, Robert E. (1978). “On the Size Distribution of Business Firms”. In: *The Bell Journal of Economics* 9.2, pp. 508–523.
- Nelson, Richard R. and Edmund S. Phelps (1966). “Investment in Humans, Technological Diffusion, and Economic Growth”. In: *The American Economic Review* 56.1/2, pp. 69–75.
- Poschke, Markus (July 2018). “The Firm Size Distribution across Countries and Skill-Biased Change in Entrepreneurial Technology”. In: *American Economic Journal: Macroeconomics* 10.3, pp. 1–41.
- Quadrini, Vincenzo (2000). “Entrepreneurship, Saving, and Social Mobility”. In: *Review of Economic Dynamics* 3.1, pp. 1–40.
- Saez, Emmanuel and Gabriel Zucman (Feb. 2016). “Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data”. In: *The Quarterly Journal of Economics* 131.2, pp. 519–578.
- Salgado, Sergio (2019). *Technical Change and Entrepreneurship*. 2019 Meeting Papers 634. Society for Economic Dynamics.
- Smith, Matthew et al. (Nov. 2019). “Capitalists in the Twenty-First Century*”. In: *The Quarterly Journal of Economics* 134 (4), pp. 1675–1745.
- Solow, Robert M. (Aug. 1957). “Technical Change and the Aggregate Production Function”. In: *The Review of Economics and Statistics* 39 (3), p. 312.
- Violante, Giovanni L. (2008). “Skill-Biased Technical Change”. In: Palgrave Macmillan UK, pp. 1–6.