

I Introduction

I.1 MOTIVATION FOR THIS BOOK

“A priority not reflected in the budget is pure demagogy”. This saying has gained popularity in the past decade among economists, political analysts, and mainstream media outlets in Mexico.¹ This expression reflects not only the political and economic circumstances of a particular country but fits the reality of budgetary decisions made by many governments worldwide. Examples of these practices abound, either for populist reasons or for lacking knowledge of how budgetary decisions affect the economy’s performance. When incoherent spending plans emerge in democratic societies, they tend to elicit a strong negative response from opposition parties or society at large. Unfortunately, budgetary proposals of this sort are not rare, so countries and their societies have to endure disappointing results down the road.

Today, it is not difficult to find cases of societal dissatisfaction with a government’s budgetary proposal. For instance, in recent years, a government in the UK collapsed in less than 50 days due to a disastrous budgetary proposal. The so-called mini-budget, as termed by the media, did not sit well with a large sector of the population, financial markets, and international organisations since it lacked a sound commitment to achieve a set of promised goals (BBC News, 2022). In another case, the minister for agriculture of India promised in 2021 to raise the target of agricultural credits. Yet, when presenting the national budget, the potential beneficiaries did not perceive any

¹ The phrase stems from an older aphorism commonly used in the Mexican political parlance between the 1960s and 1990s: *A friendship not registered in the payroll is pure demagogy*. This phrase was popularised by the satirical writer Carlos Monsiváis. Both versions refer to the asymmetry between stated intentions and real financial commitments.

real intent to improve their livelihoods. Protest leader Kirankumar Visa declared: “Forget about the targets [...] There is not even one measure to either raise farmers’ income or generate jobs” (Bhardwaj, 2021).

While expenditure dissatisfaction may not always be a society-wide movement, specific sectors – often represented by lobbying groups – frequently attempt to reflect their agenda in the government’s budgetary allocation. For example, in early 2022, the South African Federation of Trade Unions (with 800,000 members) mobilised protests in Pretoria with the aim of affecting the national budget. Its general secretary, Zwelinzima Vavi, demanded the government not only a set of policies to hike workers’ income but also to commit more resources to service delivery, schools, hospitals and police stations (Business Tech, 2022). Thus, there was a clear request to modify policy priorities through government expenditure across a large set of issues.

Perhaps, when writing this book, the most mediatic example of social protest to modify governments’ priorities is the Just Stop Oil movement. These protests have taken place in several European countries and became highly visible in mainstream media due to polemic actions such as defacing world-renowned paintings and closing motorways. Independently of whether this movement is right or wrong, their motivation has a clear budgetary target: eliminating subsidies and tax breaks for new fossil fuel extraction (Lu, 2022). Thus, these movements justify their actions by the need to make their voice heard against demagogic politicians claiming that climate change is a high priority when, in fact, they do not tackle this issue in their national budgets.

Like these examples, we can find similar others around the world, in both rich and poor countries, and across different development issues. Overall, they commonly exhibit three related features. First, they have a set of (development) goals in mind. Often, these are goals expressed by the government in official documents or campaign promises that a large societal sector agrees upon. Second, there are

specific actions by the government to achieve those goals, often materialised in terms of a budget. Protests like the ones mentioned above emerge when these expenditure actions do not seem conducive to the goals broadly agreed upon. Thus, there is a perception of policy incoherence as the state's priorities are not reflected in the budget. Third, as a response to this incoherence, lobbying groups exert pressure on governments to align their stated priorities with the budget and, sometimes, push their own agendas into the governments' priorities. That being said, different societal actors can read, from government budgets, how seriously committed a government is to achieving its development goals or campaign promises. However, such reading is not always straightforward.

While priorities on big issues such as climate change and agricultural subventions may be easily identifiable by stakeholders and the media, there exist many other development dimensions where such identification is not possible if they do not receive enough media attention. For instance, budgets (or a large part of them) may not be disclosed truthfully by governments; expenditure data may be too aggregate to make a clear connection between the expenditure and specific policy issues; certain policy interventions are inherently more expensive, which is not always easy to observe from just looking at data; expenditure patterns may carry an inertial or historical component that does not reflect the current government's priorities; and there may be fiscal, political, and bureaucratic rigidities limiting the flexibility and scope with which a government can reallocate resources. Moreover, the opacity in government expenditure data has been historically a political tool for governments attempting to obfuscate their true priorities, as unclear data limit society's auditing capabilities.

Nevertheless, in the last decades, there has been a shift towards fiscal transparency, at least in democratic societies. With the aim of making budgetary data more accessible, multiple international initiatives have emerged to strengthen the growing 'Open Government' agenda. Hence, today, we can find large

repositories of open-spending data for national and subnational governments, with various degrees of quality and resolution. Despite these remarkable efforts, the challenge of understanding policy prioritisation remains daunting since it is not just about reading true intentions through budgets but also about how such expenditure translates into development. Thus, budgetary transparency is only the first step towards a bigger challenge: *understanding the expenditure–development relationship*. This inquiry is precisely our motivation for elaborating an analytic computational framework, applying it to answer diverse questions related to sustainable development, and writing this book.

I.2 CUTTING-EDGE METHODS FOR CHALLENGING GOALS

Let us switch from the expenditure side to the development end. Here, we can say that efforts to quantify societal progress predate and are more notorious than the advances in fiscal transparency. Nowadays, many development indicators support evidence-based policymaking on specific issues. Historically, different communities of consultants and academics use specific indicators, in isolation from each other, according to their domain of expertise. In the last decades, however, the development community has shifted towards integrating several domains into a multidimensional and complexity perspective. The leading initiative is the United Nations 2030 Agenda of the Sustainable Development Goals (SDGs), which has catalysed the construction of more development indicators across a broader spectrum of policy issues (e.g., poverty, inequality, social inclusion, environmental sustainability, and public governance). Several academics and consultants have jumped into the SDG bandwagon, giving place to numerous studies analysing interrelationships between SDGs.

While these efforts have been instrumental in shifting the overall discussion of development towards a systemic framework, they remain disconnected from budgets and, hence, from policy prioritisation. That is so because they tend to focus on the ‘output’ side of the

expenditure–development relationship. Without understanding the role of government spending (the input), insights derived exclusively from indicator-based studies are somehow limited, especially when advising governments on the prioritisation of development issues. One would imagine that understanding the expenditure–development link more explicitly and the policy prioritisation process would be a primary task in an agenda that seeks to persuade governments to commit to budgetary actions in the pursuit of global goals. Unfortunately, this is not yet the case, at least not when it comes down to quantitative frameworks that operate in large multidimensional settings.

From our experience, we have come to realise that there is a lack of analytic tools to address policy prioritisation (from a systemic point of view). Because of that, efforts in promoting data transparency seem wasted. In general, governments do not get serious about promoting policy evaluations for measuring the impact of budgetary allocations on development outcomes. More specifically, there are five popular practices among governments and consultants that, in our opinion, need to be addressed to make substantial progress in policy coherence.

First, the process through which governments arrive at their development strategies (reflected in documents such as development plans, agendas to promote industrial transformation, and campaign platforms) tends to be, at best, *ad hoc* and, in the worst cases, completely arbitrary. Second, even if development plans result from legitimate democratic or professionalised processes such as national consultations and diagnostic frameworks, the data indicate that budgets rarely reflect the declared priorities. Third, traditional evaluation tools do not address systemic problems. In particular, studies employing techniques such as micro-econometrics and field experiments need to be extremely narrow to produce causal statements with the available data. Fourth, even with good data, conventional studies do not consider the myriad political-economy factors that mediate the expenditure–development relationship (such as inefficiencies and corruption). Not accounting for the political economy precludes an

accurate mapping between budgets and outcomes. Fifth, to make things even more difficult, such political-economy mediation often takes place at various scales (micro, meso, macro), so the use of aggregate data in regression analyses is insufficient to provide convincing advice.

Oddly, when we started this research programme, mainstream economists and political scientists in academia did not seem interested in tackling these issues. In development economics, a social norm, where the gold standard was to perform micro-level studies of causal impacts, had emerged. In political science, the most relevant research used aggregate expenditure data to estimate the distribution of budgetary changes. Unfortunately, the latter studies showed no intention of establishing a connection with the outcomes generated by expenditure changes. Overall, a systemic focus seemed to be irrelevant in the minds of conventional analysts, as if complications associated with inadequate policy prioritisation did not exist. For us, the data, the published research, the methods employed, the proposed solutions, and the policy practices indicated quite the opposite.²

Outside economics and political science, different research communities have long been interested in problems related to a systemic view of development. However, they did not usually share the right combination of tools and theories. For example, development studies journals started publishing some analyses using network metrics and complexity ideas, but their usual readership lacked the skills to properly embrace cutting-edge ideas coming from network and complexity sciences. In contrast sustainability scholars and network scientists were already investigating networks

² As scholars, we are aware that it is common to observe conformity with dominant ideas and a tendency to dismiss proposals attempting to defy the status quo. For example, complexity economics has been subjected to such dismissals for more than three decades (Arthur, 2021). Nevertheless, today, many ideas and methods stemming from this community have transformed practices in key economic affairs such as financial regulation (through the measurement of systemic risk), industrial strategy (by better understanding firm dynamics), and international trade (by quantifying the economic sophistication of nations).

of SDGs using state-of-the-art methods but without incorporating the critical expenditure–development connection, or its political-economy elements. Other fields that have historically focused on systems thinking (such as system dynamics) were also studying SDGs but their tools could only work at the macro level, missing causal mechanisms involving the political-economy process.

Both of us had worked for some time in the intersection between economics and complexity science. Hence, we were familiar with the methodological gaps between intriguing problems raised by social scientists and their unawareness of cutting-edge methods. Thus, it became clear that, if we wanted to push the boundaries of understanding the expenditure–development link, we had to take on some of these challenges and formulate a new analytical framework that would combine ideas and methods from various disciplines. Importantly, due to the urgency of providing well-grounded advice for closing the development gaps of the 2030 Agenda, any proposal would have to be empirically usable, easy to scale, and flexible enough to accommodate budgetary data with different levels of granularity. Furthermore, our framework would need to become a contributing factor in the capacity-building of policymakers and their analysts.

I.3 THE 'POLICY PRIORITY INFERENCE' RESEARCH PROGRAMME

In 2018, and in collaboration with Florian Juárez-Chávez, we published a paper entitled "*How do governments determine policy priorities? Studying development strategies through spillover networks*" (Castañeda et al., 2018). This work was our first attempt to discern the expenditure–development relationship from a systemic point of view. There, we opted for a computational approach because it allowed us to explicitly model some critical political-economy mechanisms that mediate the expenditure–development link while, at the same time, allowing us to work in a multidimensional setting. On the empirical front, we discarded using statistical methods, such as regression analysis and machine learning, since they were inadequate when

working with many development indicators (perhaps hundreds) and relatively short time series of development indicators.

Accordingly, we had to develop a framework in which a theoretically rich computational model would supplement the lack of data. This approach would allow us to generate synthetic series that preserved essential features of the empirical data, bypassing the need for gigantic development datasets and enabling scalability. The lack of disaggregated budgetary information for many countries meant that we had to model how a government allocated resources in a system where bureaucrats mediated their use and where development indicators impacted each other through spillovers. For this reason, precisely, we decided to refer to our research programme as 'Policy Priority Inference' (PPI).

Following this paper, we kept enriching our model, accessing new datasets, developing better calibration methods, and tackling more specific empirical challenges such as quantifying policy resilience (Castañeda and Guerrero, 2018), measuring policy coherence (Guerrero and Castañeda, 2021b), and estimating the effectiveness of the rule of law (Guerrero and Castañeda, 2021a). While starting as a typical research programme within the boundaries of academia, it suddenly turned into an ambitious agenda with policy implications. This take-off occurred when members of the Bureau for Latin America & the Caribbean of the United Nations Development Programme (UNDP) read our first publication and invited us to their New York offices to give a workshop. Around the same time that we were working on our first PPI paper, the United Nations had given the UNDP the mandate of coordinating much of the activities related to the 2030 Agenda. Such coordination involved facilitating analytic tools that governments could use to inform policy prioritisation towards the SDGs. Thus, our initial work seemed relevant to the UNDP's new tasks.

We quickly realised that the need for new analytic frameworks was more pressing than we originally had imagined. It was clear, from these interactions, that the models and empirical methods provided,

at that time, by economists and other communities were insufficient, as the SDGs suddenly shifted the discussions to the domain of complex systems. Our adoption of a non-conventional approach implied that we had many methodological challenges ahead, conceptual, technical, and data-related. The evolution of PPI resulted from cross-fertilisation between our academic work and our engagement in policy projects. In collaboration with the UNDP, we conducted more than ten projects where we had to adapt PPI to address highly specific empirical questions that local and national governments wanted to address.

Recurrent questions inspired by real policy problems created the need to make more realistic assumptions, which is not straightforward with more conventional frameworks. Some examples of these assumptions include the differentiation between indicators where a government has relevant intervention programmes and where it does not; allowing for fiscal rigidities that limit the government's ability to reallocate funds; enabling systemic impacts for the identification of development accelerators; and incorporating the substitutability between private and public (internal and external) sources of funding to produce advances in different development issues. Many of these elements were absent in the scholarly literature, at least in the quantitative one, so it was up to us to push the frontiers and propose ways to account for them.

During this exciting process, we received the support of different institutions that complemented the efforts of the UNDP. The Alan Turing Institute (Omar's home institution) provided, all along this process, financial means and other resources; the United Kingdom's Engineering and Physical Sciences Research Council and Economic and Social Research Council funded various parts of this research; the National Laboratory of Public Policy, at the Centre for the Research and Teaching of Economics (CIDE, Gonzalo's home institution), organised workshops with public servants, which were helpful to inform, validate, and socialise PPI; the Mexican National Council of Science and Technology funded, in part, Gonzalo's

sabbatical in the academic year 2018–2019; the Global Initiative for Fiscal Transparency brought us closer to potential beneficiaries, and the Centro de Estudios Espinosa Yglesias and the International Monetary Fund opened new dissemination spaces.

This book is the synthesis of the PPI research programme. It provides a computational framework for studying sustainable development and policy prioritisation in a multidimensional and complex setting. In it, we present the latest version of PPI and its theoretical underpinnings. We walk the reader through various empirical applications using datasets on development indicators and government expenditure with different levels of granularity. While much of the content builds on our academic publications, this book is not a collection of papers but an integrated guide through a novel methodology. All chapters are interconnected, sometimes by data, others by the simulation strategy, and others by the scope of the study. In these chapters, we employ the same model and the latest version of the data available when we started working on this book. In trying to reach the broadest possible audience, we aim to achieve a healthy balance between technical level and policy readability. Therefore, the book touches only on the technical elements needed to understand the analysis, leaving more sophisticated details to the academic publications on which it builds.

In this book, we do not intend to make a comprehensive review of other existing frameworks or studies for analysing SDGs,³ but only to provide a straightforward narrative encompassing challenges in the study of sustainable development and to present our analytic tools to tackle them. There are five main contributions of this book. First, it provides a comprehensive and flexible quantitative framework that uses modern computational methods and the most recent data available. Second, it presents a collection of empirical works tackling essential inquiries from the sustainable development literature. Third, it is a technical and detailed example of how computational

³ The reader may refer to our scholarly work for that purpose.

models can support social scientists and development consultants in approaching complex problems. Fourth, it facilitates an empirically relevant introduction to computational social sciences. Fifth, it serves as a manual or tutorial on the PPI framework.

Throughout the book, the reader is familiarised with diverse methods; for example, agent computing, gradient-search optimisation, Gaussian processes, and network analysis. However, agent computing is the principal method employed for developing the PPI framework. Agent computing – or agent-computing modelling – is an artificial intelligence (AI) framework known by other names across various fields.⁴ A popular name to refer to agent computing is agent-based modelling. However, the term agent-based modelling is sometimes confused with neoclassical rational models from economics. The reason is that, before digital computers, economists were already elaborating agent-level rational-equilibrium models. Hence, today, some economists that are not familiarised with a complexity view (especially in the public sector) tend to understand the term 'agent-based' as rational agents in an equilibrium world, which, ironically, are assumptions that the agent-computing literature leaves aside. Axtell (1999) proposed the term agent computing, as it reflects closer to what these models do: using the agents to compute limited information and interact in a decentralised fashion in an uncertain environment. Nevertheless, the term agent-based model is also widely accepted as a standard for agent computing. For this reason, we use agent computing and agent-based model (ABM) interchangeably throughout the book.

In disciplines such as sociology and epidemiology, agent computing is used extensively but not in others. In economics, despite a large body of literature employing agent computing (Axtell and

⁴ Sociologists often use the term "actor-oriented models", ecologists and epidemiologists call them "individual-based models", computer scientists say "multi-agent systems", and sustainable-development scholars refer to "bottom-up multi-agent models". While each of these names conveys different epistemological elements, they use similar analytic tools.

Farmer, 2022), it is still a method that, according to several scholars, does not fit the methodological canons of conventional thinking. While discussing the reasons for the marginalisation of agent computing in economics is beyond the scope of this book, it is worth mentioning that a common argument is the lack of reliable/robust empirical applications. We disagree with this argument, especially since the past decade has seen tremendous progress in using this framework for empirical studies. Notwithstanding, we do see a lack of seminal books that walk the reader through the entire process of developing a theoretically rich model, demonstrating how to use it empirically with statistical rigour, and deriving policy implications from the resulting inferences.⁵

A key contribution of this book is precisely filling the aforementioned gap. In other words, this is not a book about a 'toy model' nor a collection of models. The book takes the reader through a series of theoretical arguments motivating the use of agent computing in the study of countries' development, followed by the construction and validation of a model of policy prioritisation. Then, it demonstrates how the model can be employed to devise different empirical strategies to deal with a diverse set of development-related problems. In addition, and in the spirit of promoting open science and transparency, the book comes with companion code and data that allow the reader to replicate and modify all the analyses and figures.⁶

I.4 TARGET AUDIENCE

The target audience of this book is quite diverse. Let us divide it into two groups: researchers and practitioners. In the research camp, three major communities gather scholars from several fields. First, there

⁵ The seminal books on the subject were written in the 1990 and consisted mainly of theoretical models. Among the most important contributions, we can find Epstein and Axtell's Sugarscape (Epstein and Axtell, 1996) and Axelrod's prisoner's dilemma tournaments (Axelrod, 1997a). Other posterior books consist of collections of papers.

⁶ The repository of data and code can be accessed at <https://github.com/oguerrer/CESD>.

is the development community with three subdivisions: development economics, development studies, and sustainable development. Each of these communities is large and includes a diverse set of scholars. The second group encompasses academicians in political science and public administration. In particular, policy change and public financial management are highly relevant fields as they study government expenditure and its societal implications. Third, it includes a community of interdisciplinary studies that spans beyond social scientists and brings together researchers from fields such as physics, computer science, and applied mathematics. Here, we can find the growing communities of complexity economics, complexity science (which includes network science), systems analysis, and computational social science. Attracting the attention of all these heterogeneous communities is not an easy endeavour. For this reason, the book strives for a balance between theory, technical detail, and empirical analyses.

In the practitioner audience, we can also find a diverse readership. First, we have governments, local and national. More specifically, treasuries, budget and planning offices, and their technical experts could greatly benefit from this tool as it can easily take advantage of much of the data they have built. Second, multilateral organisations and consultants seeking to support capacity building or to perform global assessments are also potential users. Third, NGOs and civil society at large could use PPI to evaluate the government's priorities and strategies, or to propose alternative ones. Fourth, legislative institutions could use this tool as part of the budget approval process. Fifth, political parties can employ PPI to assess the feasibility of political campaigns (of themselves or their competitors). Sixth, think tanks could benefit from integrating this framework as part of the analytic toolkit of their consultants. Seventh, international partnerships promoting open data and transparency agendas could use PPI as an example of how such data could be exploited and used for the benefit of society (i.e., for advocating better budgeting).

I.5 STRUCTURE OF THE BOOK

We organise the book into three parts. The first introduces the theoretical framework and the empirical datasets that serve as the backbone for the computational model of PPI. The second part presents applications with a global perspective. That is, we demonstrate how to apply PPI to various problems typically discussed on a comparative cross-country basis. In the third part, we descend into more disaggregate levels of analysis, both conceptually and data-wise. In these chapters, we introduce the frontier in expenditure data and demonstrate how PPI can exploit their nuanced structures. Overall, the book structure establishes, in the first stage, sound conceptual and theoretical grounds. Then takes the reader, gradually, through a sequence of increasingly more nuanced and sophisticated applications of PPI. Next, let us summarise the content of each of the remaining chapters.

In Chapter 2, we start with a brief exposition of how PPI analyses multidimensional development in a setting of interdependent policies and uncertainty. Then, we explain how to produce causal inference through agent computing. This part is highly relevant for the analyses presented throughout the book, and for contrasting the capabilities of any ABM with those of popular econometric methods. After clarifying that there are different accounts of causality in the literature, we argue that agent computing combines data and theory (social mechanisms that connect the micro and macro levels) to produce artificial data-generating processes. Through simulations, the method provides counterfactuals suitable for assessing the performance of a system that exhibits a policy intervention. Lastly, we mention ten advantages that agent-computing models offer for measuring causal impact in contrast to econometric approaches.

In Chapter 3, we introduce a popular dataset on development indicators and aggregated data on government expenditure across a large set of countries. These two types of data constitute a comprehensive picture of the current state of development across multiple dimensions and geographies, as well as our first glimpse of

governments' heterogeneous budgeting capabilities for improving these indicators. Then, we review some of the most popular methodological frameworks to analyse these data and to try to understand the expenditure–development relationship. We highlight their limitations to explain why there has been little progress in studying development from a multidimensional perspective and even less in estimating expenditure impacts from a systemic point of view. This presentation sets the ground for a more thorough discussion of what we see as empirical challenges that need addressing to make progress, both in research and policy.

Taking the aforementioned challenges as motivation for a new framework, Chapter 4 presents the PPI model in full detail. We elaborate further on data considerations and the empirical features or regularities that our model tries to reproduce. Then, we describe causal mechanisms binding government spending, political economy, and development indicators together. Along the way, we elaborate on the socioeconomic intuition of each component. Finally, we present an overview of the model structure through its parameter space, a structural diagram, and pseudo-code describing its dynamics.

Chapter 5 is for the more technically knowledgeable reader. It is not necessary to go over it to understand the applications of PPI; nonetheless, a reading is convenient if one has more specific inquiries on how to make empirically relevant agent-computing models. It covers various aspects regarding the indirect calibration of parameters for cases lacking data for direct calibration. Alongside the development of a gradient-search calibration algorithm, we discuss important aspects of statistical inference from computational models, such as confidence intervals and hypothesis testing. Then, we elaborate on the various validation aspects of agent computing. Finally, we discuss the virtues of this type of tool in enabling counterfactual analysis and, hence, causal inference.

We begin Part II with Chapter 6, which presents the first application of PPI. In this chapter, we perform a prospective analysis of the potential trajectories of various development indicators across a large

sample of countries. The analysis estimates development gaps that would remain open by 2030 should similar government expenditure levels remain in place. Then, we perform counterfactual simulations that vary the expenditure usage and estimate the sensitivity of the different development dimensions to budgetary changes in a large set of countries.

Next, in Chapter 7, we build on these estimates and introduce a new concept: *the budgetary frontier*. This idea consists in simulating the effects of unrestricted government expenditure (the frontier) and studying the performance of the indicators under such a setting. If some gaps remain open by 2030 while operating on the budgetary frontier, this result indicates the potential existence of idiosyncratic structural bottlenecks. In other words, we infer the presence of structural factors – unrelated to government expenditure – that dampen development and render government spending ineffective. We call them idiosyncratic bottlenecks because they are specific to individual policy issues (in Chapter 11 we analyse another type of bottleneck). Finally, we develop a flagging system to pinpoint the potential presence of these bottlenecks and differentiate whether poor performance is due to underfunded policies or idiosyncratic bottlenecks.

In Chapter 8, we maintain a global perspective but shift to a more specific type of expenditure: the one dedicated to improving the implementation of public governance programmes. Political economy is a central element of PPI, and it uses data about the quality of public governance to inform some of its parameters. In this chapter, we investigate how public expenditure can impact the efficiency of government programmes when those parameters are part of the endogenous development indicators. We perform this study by developing a novel conceptualisation of the expenditure channels that can affect public governance, distinguishing between relative and total changes in the relevant spending. We also devise a simulation strategy to characterise the policy landscape (across the different expenditure channels) governments face when trying to reform public governance. Roughness in this landscape quantifies the uncertainty of potential

outcomes of governance-related expenditure. Thus, countries with a rugged policy landscape will face a more difficult path towards institutional development.

The global perspective conveyed in Part II ends with Chapter 9 and the study of aid effectiveness. This study pushes the boundaries of the aid-effectiveness literature, where international assistance has been historically treated, in quantitative models, as a black box. Here, we exploit PPI's flexibility to enrich the aggregate expenditure data of countries by including granular incoming aid flows in a large sample of developing countries. These data contain information on millions of aid projects categorised into the SDGs, allowing us to analyse the aid-indicator relationship. The effectiveness of these flows is assessed by devising counterfactuals in which aid is absent and by developing a new impact metric that accounts for non-linear changes in the performance of indicators. We employ this impact metric in the subsequent chapters of the book. It is worth noting that this chapter provides the first large-scale study of aid effectiveness estimated at the level of each country and indicator.

In Part III, we shift from a global to a focalised perspective. This extension means that we implement the model in a single country and specific development issues. Our country of choice is Mexico, as it is a leading nation in the publication of open spending data with a high level of disaggregation, temporal coverage, and SDG interlinkages. In Chapter 10, we study the problem of subnational development. More specifically, we look at the challenge of financing multidimensional development across 32 states that exhibit high fiscal imbalances due to their limited tax-collection capabilities and substantial expenditure needs. We assemble a unique dataset with more than 100 state-level SDG indicators and state-level budgetary data. Then, we study fiscal federalism by assessing how Mexico can establish different cross-state budgetary distributions to achieve shared development goals.

Next, in Chapter 11, we investigate a concept that has become very popular in the SDG policy community: *development accelera-*

tors. Conceptually, development accelerators are policy issues that, if financially supported, can catalyse development over a broader set of topics through spillover effects. Hence, by definition, they are systemic. To the extent of our knowledge, there are no solid quantitative frameworks to identify accelerators or analytically well-grounded estimations of their systemic impacts. This chapter provides the first quantitative method to achieve this. We design a protocol to estimate systemic effects through randomised simulated financial interventions. Since we develop the notion of systemic effects, we introduce and analyse a second type of bottleneck: *systemic bottlenecks*. In contrast to idiosyncratic bottlenecks, systemic ones are the counterpart of accelerators. If not financially supported, they become anchors that slow down the development of other policy issues. In this study, we work with a novel dataset of public expenditure in Mexico linked to the 169 targets of the SDGs. This chapter provides the first rigorous study of development accelerators and systemic bottlenecks.

Chapter 12 presents the last application of PPI. Here, we study the problem of socioeconomic deprivation, which is multidimensional, complex, and central to the Mexican government and society (as well as to most of the world). We assemble a large-scale dataset of specific expenditure programmes related to socioeconomic development (about 700 programmes), each one linked to at least one development indicator. Thus, this chapter takes full advantage of the capabilities of PPI to handle highly structured datasets. We also demonstrate its flexibility by modifying one of the key equations of the model to take into account private sources of funding (i.e., household income) that are complementary to government expenditure in achieving social advances. This extension allows us to estimate the impact that individual government programmes, household income, and US remittances have on socioeconomic-deprivation indicators. It also facilitates studying the government's capacity to mitigate the adverse impacts of income drops through its various programmes (and their effectiveness). This chapter provides the reader with the state

of the art in ABM's impact evaluation techniques and government spending data.

Chapter 13 is the final chapter. First, we summarise the lessons learnt through the various analyses presented in Parts II and III. Then, we provide the reader with a workflow to elaborate policy guidelines from the results of our simulation exercises. This workflow is a systematic way to "make sense" of the different kinds of analyses presented throughout the book and to guide policymakers in their interpretations. This scheme of policy guidelines evolved from the experience gathered through our various policy-oriented projects. Despite the usefulness of these guidelines to policymakers, they are not usually part of the toolkits of technically inclined analysts. Finally, we close the chapter and the book with a reflection on the skill set that social scientists need to acquire if they want to exploit computational frameworks such as PPI, and advocate for further investments in developing the emerging field of Computational Social Science.