

(When) Do Open Budgets Transform Lives?

Progress and Next Steps in Fiscal Openness Research

By Martin Haus, Joachim Wehner, and Paolo de Renzio



We document a rapidly growing empirical literature that can plausibly claim to identify causal effects of transparency or participation in budgeting in a variety of contexts. Recent studies convincingly demonstrate that the power of audits travels well beyond the context of initial field-defining studies, consider participatory budgeting beyond Brazil, where such practices were pioneered, and examine previously neglected outcomes, notably revenues and procurement. Overall, the study of the impacts of fiscal openness has become richer and more nuanced. The most well-documented causal effects are positive: lower corruption and enhanced accountability at the ballot box. Moreover, these impacts have been shown to apply across different settings. We conclude that the empirical case for open government in this policy area is rapidly growing in strength. We set out challenges related to studying national-level reforms; working directly with governments; evaluating systems as opposed to programs; clarifying the relationship between transparency and participation; and understanding trade-offs for reforms in this area.

Open and accountable government budgeting remains an essential component of good governance. Most recently, the need for open budgeting has been thrown into sharp relief across the world by the large-scale emergency expenditures in response to the Covid-19 pandemic. Yet proponents of greater fiscal openness are challenged to provide credible evidence that transparency and participation can deliver real benefits to those who are most reliant on government services, and society more widely. One comprehensive review of evidence on the impacts of fiscal openness considered 38 empirical studies published between 1991 and early 2015 (de Renzio and Wehner 2017). It found that increased budgetary disclosure and participation are consistently associated with improvements in the quality of the budget, as well as governance and development outcomes. At the same time, this review also pointed out that only a handful of studies convincingly identify causal effects, in the form of reduced corruption, enhanced electoral accountability, and improved allocation of resources. Causal evidence was especially thin for ultimate impacts of interest related to human development, for example health outcomes such as infant mortality, or educational attainment.

Given how much the field has evolved and the growing interest in this research area—and in open government more generally—it is both opportune and necessary to update the review with more recent findings from published academic research and other rigorous policy-oriented research, covering the years from 2015 onwards. We also want to reflect more broadly on the state of research on the impacts of fiscal openness, to shape the debate on its future directions. Our research was commissioned by the International Budget Partnership (IBP), which works in collaboration with civil society, state actors, international institutions, and the private sector to create open and accountable government budgeting processes, and by the Global Initiative for Fiscal Transparency (GIFT) network, which helps countries to advance and institutionalize significant, continuous improvements in fiscal transparency, participation, and accountability. It also contributes to a review by the Open Government Partnership (OGP) taking stock of available evidence on several topics relevant to open government, including fiscal openness.

The immense interest in understanding the impacts of transparency and participation more generally is reflected in a range of related reviews that have been undertaken since 2015, including by Alt (2019), Khemani et al.(2016), Kosec and Wantchekon (2020), OGP (2018), Rathinam et al.(2019), Rudiger (2018), and Tsai et al.(2019). A recent review by the Abdul Latif Jameel Poverty Action Lab (J-PAL 2020) considers a crucial segment of the work within the scope of our review. While broadly related to these other efforts, our survey is much more specific in scope, as we set out further in the following section. We focus exclusively on the impacts of government openness in the budget process.

We are interested in understanding and documenting what happens when governments open up the way they manage public finances in practice, as debates about the merits of open budgets cannot be settled on the basis of normative arguments alone. Fiscal openness enthusiasts should test their beliefs and expectations, not only in search of supportive findings, but also to understand potential limitations and how these might be overcome. Skeptics, on the other hand, will only be swayed by convincing empirical evidence. Our focus also connects directly with the earlier survey by de Renzio and Wehner (2017).

In the following, we set out the scope for this survey and outline our approach to identifying relevant literature. Next, we highlight key findings and several subsets of the literature more specifically. The conclusion relates the discussion to gaps and patterns identified in the earlier review and develops implications and recommendations for further research.

1. Review scope and process

This review covers research published since 2015 that empirically evaluates a causal argument about the impact of fiscal or budget openness. The term “fiscal openness” or “budget openness” encompasses both transparency and participation in relation to the budget, including spending and revenues, as well as deficits and debt. More specifically, it refers to government action that involves the official publication of budget information (transparency) or public authorities providing opportunities for citizen engagement (participation) at any stage of the budget process. The budget cycle covers the formulation, approval, execution, and audit and evaluation stages. Our review thus excludes purely theoretical treatments without empirical evidence, or empirical work examining interventions undertaken by actors other than the government, such as civil society organizations, academic researchers, private companies, etc.

Principally, we looked for research that achieved publication as a peer-reviewed academic article, a book with an academic press or well-known commercial publisher, or as a working paper of an institution that has made significant contributions to work in this area. We did however make some exceptions for manuscripts in development and unpublished at the time of our review (the autumn of 2021), which in our assessment revealed significant promise and cover substantively important aspects in this area. Moreover, we did not constrain our review in terms of the types of impacts, which could be any outcome broadly related to human or economic development, governance, or the use of public financial resources in aggregate, allocative, or operational terms (see de Renzio and Wehner 2017: Figure 1). This also includes equity, for example in relation to who participates or benefits from fiscal openness.

Yet we applied the above rules flexibly and remained open to considering closely related papers that have direct implications for understanding the impacts of fiscal openness. This might be where an intervention was undertaken by actors other than the government but had a clear link to transparency or participation in public finance, for instance when it was implemented in close collaboration with government authorities, or at substantial scale. Also of interest were conceptual critiques and discussions, studies illuminating important contextual conditions for example related to civil society or the media, or methodological approaches that might be deployed in future research on the impacts of fiscal openness.

A final criterion was the strength of the underlying research design. Given our search for evidence on impacts, our focus here is on studies that can credibly claim to identify causal effects, especially those based on randomized evaluations (Glennerster and Takavarasha 2013) or natural experiments with (as-if) random treatment assignment and regression discontinuity designs (Dunning 2012). Field experiments are widely considered the gold standard for causal inference, but they can be difficult or impossible to implement, for instance to study economic policy. Natural experiments, too, may not be available to examine a context or phenomenon of interest. Hence, we also included studies that are not purely correlational but fall short of being fully experimental. These limit some threats to valid inference, for example by using fixed effects estimators, or deploy quasi-experimental approaches with non-random treatment assignment, such as differences-in-differences (Angrist and Pischke 2009).

For the survey underpinning this review we emailed 55 individuals across universities, think-tanks, and multinational organizations. The respondents were scholars or practitioners with a significant track record of work in this area of study. We consulted colleagues at IBP and GIFT in compiling this list. Some of the recipients forwarded our request to colleagues and responded with joint input. Altogether, we received more than 30 responses of varying length, including several very detailed ones. Our survey covered three questions that asked respondents about their main contributions related to this area, what they regarded as major contributions by other scholars, and their views on future directions for research in this area (see Appendix A).

We collected all studies recommended by the survey respondents and added some that we had identified ourselves. This resulted in a total of more than 200 publications or papers. In a next step, we applied our scope criteria to this entire set of publications. We thus looked for studies (a) published or produced since 2015 that (b) had an independent variable or “treatment” that related to government disclosing information or enabling participation at any stage in the budget process, (c) analyzed this in relation to outcomes including in the broad areas of governance, development, or the quality of the budget, and (d) could make a claim to identify causal effects or present robust statistical associations based on research designs that protect against some threats to valid inference, principally those discussed by Angrist and Pischke (2009).

Two of us individually went through the papers and applied the above four criteria, and then jointly agreed our selection. This yielded a total of 32 studies, summarized in Appendix B, which complements the appendix in de Renzio and Wehner (2017) listing the studies included in the earlier review. Appendix B indicates, amongst others, the main dependent and independent variables, and whether a study was based on a randomized controlled trial, natural experiment, or quasi-experiment. We also collected some additional research that did not comply with our criteria, but with a potential to inform conclusions and recommendations in our paper, such as the above-mentioned surveys of the literature in overlapping or adjacent fields. The references at the end of this paper distinguish studies that were included in our review and others that we refer to in the wider discussion.

2. Overview of findings

We detect substantial scholarly interest in both fiscal transparency and participation since the previous review was carried out in 2015. Our list includes 30 studies covering effects of either participation (14) or transparency (16), and two that examine both within the same research design. It is striking that the strongest research designs are predominantly found in studies relating to the impacts of budget transparency, especially of audits. Of the 12 studies that can make convincing claims to identify causal effects due to their research designs, nine involve audits, one relates to participation, and two examine both audits and participation in the same research design. Three of these 12 studies do not meet all our criteria in the strictest sense, as they do not directly assess the impact of a government openness (Beuermann and Amelina 2018, Keefer and Roseth 2021, Sjoberg et al. 2019). However, we decided to include them as they provide a rare instance of a randomized controlled trial covering participation (Beuermann and Amelina 2018), focus on service delivery outcomes rather than more immediate impacts (Keefer and Roseth 2021), and provide cross-country evidence on revenues (Sjoberg et al. 2019). Before discussing the main substantive developments, this section first highlights some wider patterns.

One noteworthy development is that scholars of participatory budgeting have shifted attention to new contexts. In addition to further work on Brazil (Touchton et al. 2020, Touchton and Wampler 2020) and India (Costa-Font and Parmar 2021), recent contributions include evidence from Indonesia (Grillos 2017), the Philippines (Saguin 2018), Russia (Beuermann and Amelina 2018), South Korea (Hong 2018, Jung 2021, No and Hsueh 2020) and the United States (Calabrese et al. 2020, Hagelskamp et al. 2020, Karner et al. 2019, Shybalkina and Bifulco 2019, Su 2018). Most of the latter—four out of five studies cited—focus on participatory budgeting in New York City. This focus reflects the city’s density of academic institutions and social scientists, which raises an important point related to the selection of cases that get documented in the literature. While the overall diversification of cases is very welcome, funders and the research community should pay attention to ensuring adequate coverage of cases that are less accessible or where local scholars lack prerequisite resources.

Our summary also reveals distinct patterns of academic attention across different disciplines in the social sciences. Studies of participation in budgeting are dominated by public policy scholars and political scientists, while work on transparency, including audits, is more dominated by (political) economists. This may be related to different substantive focuses and methodological approaches. For example, political scientists may be more interested in the nature and quality of deliberative processes and deploy qualitative research as well as quantitative research designs. It is unclear whether the dearth of causally identified evidence on the impacts of participatory budgeting reflects disciplinary differences in research traditions, or an inherent difficulty to identify impacts in this area. In contrast, all four studies that, at the time of this review, had achieved publication in major general or field journals in economics examined impacts of audits (Avis et al. 2018, Bobonis et al. 2016, Larreguy et al. 2020, Zamboni and Litschig 2018).

In terms of substantive progress, it is nice to see that several gaps identified previously by de Renzio and Wehner (2017) are getting more attention in this recent wave of research. Examples include procurement (Bauhr et al. 2019, Geradino et al. 2017, Zamboni and Litschig 2018) and revenues (Sjoberg et al. 2019, Touchton et al. 2020).

In addition, there is new research on traditional macro-fiscal outcomes that have been linked to transparency, including foreign direct investment (Cicatiello et al. 2021) and access to borrowing (Kemoe and Zhan 2018). In the wake of the global financial crisis in 2007 and 2008, increased attention has been given to assessing and ensuring the disclosure of fiscal risks, as reflected in the IMF's revised Fiscal Transparency Code. Recent scholarship picks up on this important new focus. ElBerry and Goeminee (2021) report a correlation between fiscal risk monitoring and budget credibility in low and middle-income countries. Copelovitch et al. (2018) contribute a study that links government disclosure of financial system data to lower and less volatile borrowing costs in advanced economies with low debt burdens. Aaskoven (2020) shows that access to revenues from oil amplifies electoral cycles in public spending, but that this can be countered by budget transparency. These are substantively important aspects to examine, even if it is difficult to do so with causally identified designs. Further work on transparency and macro-fiscal outcomes will benefit from more studies with designs that mitigate some endogeneity concerns or in settings that offer additional causal leverage (e.g., Glennerster and Shin 2008).

Lastly, we note several studies tackle issues that may lead to a more nuanced understanding of the potential of open budgeting. Studies from different contexts provide evidence and caution that participatory processes with bias in representation may not effectively target those most in need (Grillos 2017, Karner et al. 2019, Saguin 2018). Bauhr and Carlitz (2020) present data from Vietnam suggesting that whether transparency improves public service delivery may depend on the nature of the service, with improvements more likely where service provision entails substantial discretion by street-level bureaucrats. Some of the literature on audits looks at potential side-effects, such as distortions of procurement procedures (Gerardino et al. 2017), or the crowding-out of grassroots participation in bottom-up monitoring (Gonzalez et al. 2020). This is not to detract from an overall positive policy message in this literature. Rather, such studies are a sign of a maturing research field where attention is shifting from establishing whether transparency and participation “matter”, to probing the conditions that enable benefits to be realized, and potential trade-offs and unintended consequences. In more technical terms, this allows for a more nuanced assessment of impacts beyond the Average Treatment Effect (ATE) towards including moderating variables that may not only matter for the impact size, but at times also determine whether the treatment has a positive or negative effect.

3. Specific results

In this section, we focus on four substantive areas where our review identifies significant activity or progress. The first two examine specific independent variables. Here, we review studies documenting the effects of external audits, an area where the evidence is most convincing due to several studies with natural experimental designs that involve (as-if) random treatments, and participatory budgeting, which is now studied across many more empirical settings than Brazil. In addition, we also discuss work where the dependent variable is related to revenues or procurement, where scholars have started to fill research gaps noted in the previous review. While this implies some overlap of the studies covered in the first two and the last two sections, it allows us to highlight the main developments. This presentation should also benefit readers interested in any one of these four areas.

Audits

The most extensive documentation of causal impacts within the scope of this review relates to new research on external audits. All of the studies discussed here involve external audits whose results are made publicly available. Earlier seminal contributions documented audits reduced the extent of corruption in Indonesian road projects (Olken 2007) and facilitated electoral accountability of Brazilian mayors (Ferraz and Finan 2008). More recent work returns to these and related questions in the contexts of Indonesia (Gonzalez et al. 2020) and Brazil (Avis et al. 2018, Zamboni and Litschig 2018, Funk and Owen 2020), but also broadens the geographical scope by providing evidence from Chile (Gerardino et al. 2017), Italy (Vannutelli 2021), Mexico (Larreguy et al. 2020), Puerto Rico (Bobonis et al. 2016), and South Africa (Berliner and Wehner 2021). These studies can make causal claims due to their research designs. We distinguish randomized controlled trials (Gonzalez et al. 2020, Zamboni and Litschig 2018), natural experiments with (as-if) random treatment assignment (Avis et al. 2018, Berliner and Wehner 2021, Gerardino et al. 2017), and quasi-experiments without random treatment assignment (Bobonis et al. 2016, Larreguy et al. 2020, Vannutelli 2021).

Following the tradition of Olken (2007), one set of studies examines causal effects of audits on corruption or service delivery (see also J-PAL 2020). A reanalysis of Olken's data by Gonzalez et al. (2020), which we discuss more fully in the next section, confirms the effectiveness of audits but suggests they weaken the effectiveness of community monitoring. Zamboni and Litschig (2018) find that in Brazil, a randomized increase in the annual audit risk by 20 percentage points reduced both the share of audited resources for procurement and of procurement processes with evidence of corruption by about 10 and 15 percentage points, respectively. On the other hand, these audits do not affect satisfaction with health services or compliance with regulations for a cash transfer program, which the authors attribute to differences in the severity and likelihood of sanctions. Also on Brazil, Avis et al. (2018) build on Ferraz and Finan (2008) and find that random audits reduce corruption, predominantly by increasing the chances of a police crackdown or a conviction in court. Beyond corruption, Funk and Owen (2020) revisit the Brazilian audit program and document positive impacts of audits on an index of municipal service delivery outcomes. Bobonis et al. (2016) analyze the impact of audits on corruption across municipalities in Puerto Rico. In this setting, mayors can anticipate the timing of audits.

The study shows they adjust corrupt behavior when this could result in a damaging audit ahead of elections—only subsequently to revert to higher levels of rent extraction. This finding also underlines the need to consider long-term impacts rather than basing policy recommendations only on ex-post impact evaluations shortly after an intervention.

A second set of studies follows Ferraz and Finan (2008) to examine electoral outcomes. Bobonis et al. (2016) find that mayors who adjust to pre-election audits by reducing corruption increase their likelihood of re-election. But as these audits do not induce a sustained reduction in corruption, voters are left with more responsive but not less corrupt politicians. Larreguy et al. (2020) find the presence of local broadcasters conditions the effect of pre-election audits in Mexico. An additional local station reduces the vote share of the incumbent mayor's party by about 1 percentage point where audits detect substantial unauthorized or badly targeted expenditures, and vice versa. Berliner and Wehner (2021) examine the impact of a nationwide summary report on local audits in relation to as-if randomly timed by-elections due to the death of local ward councilors in South African municipalities. Here, the vote share of the party of the mayor responsible for municipal funds increased (decreased) in response to improved (poor) audits, by about 5 percentage points, but only when by-elections occurred shortly after the nationwide release. This effect is not conditional on the presence of local media, but they show local audits get significant media attention in general and affect how citizens evaluate the performance of their mayor and local authority. All studies highlight the need for audits to be timely—available ahead of elections—to facilitate accountability.

Others add new angles that contribute findings on when or how audits matter. Vannutelli (2021) exploits a change in Italy where the appointment of auditors was removed from the discretion of mayors and randomized instead. The reform was introduced in a staggered manner depending on the timing of the expiration of an incumbent auditor's three-year term. The study finds significant improvements in budget management due to the reform, resulting in an increase in the average net surplus of about 1 percent of the overall annual budget. The study is a reminder that the independence of auditors is important, highlighting the need to differentiate impacts along the quality dimension. This is likely a problem especially where auditors have strong local ties or dependencies, or where the quality of auditors deployed can be systematically manipulated. Another study, by Gerardino et al. (2017) highlights how auditees may strategically adjust to audits. Looking at procuring public entities in Chile, they find that audits induced a shift to less transparent direct contracting that made it more difficult to detect infractions. For audits to be fully effective, the potential for such distortions needs to be considered and countered.

Overall, this impressive subset of literature documents that the publication of independent external audits can reduce corruption and enhance electoral accountability, especially where their public dissemination is supported by local media. The effectiveness of these external audits is reduced where auditors lack independence, and where auditees can anticipate the timing of audits relative to elections, or otherwise adjust their behavior strategically to weaken the detection or consequences of poor governance.

Participation

Most recent studies of participation lack a robust causal design based on (as-if) random treatments. As an exception, Beuermann and Amelina (2018) present a randomized controlled trial of participatory budgeting in rural Russia.

This took place within the context of a broader decentralization reform that included a legal mandate for public hearings before the budget gets approved and for citizen monitoring of spending through local councils. The intervention was part of a World Bank project implemented by a non-government entity and hence was, strictly speaking, outside the scope of our review. Yet we include it, as it provides a rare instance of an experimental design in the participation literature.

As part of the intervention, local authorities were randomly assigned training and technical assistance. The evaluation showed that training alone was unable to improve the implementation of participatory budgeting, while in combination with technical assistance it improved a variety of outcomes, including citizen engagement, local tax revenue collection, preference matching between citizens and authorities, and public service satisfaction. Yet, these positive findings only materialized in “mature decentralized contexts” (Beuermann and Amelina 2018: 376), indicating that institutional history matters for enabling effective participation and echoing results from Brazil that we discuss below. Outside such a context, the provision of technical assistance was only able to increase tax revenue collection. The intervention also underlines the importance of better understanding how administrative capacity matters for effective participation. It hints at a potential trade-off in contexts where such capacity is rare (if scarce administrative resources are used for making participation effective, does this undermine other activities one might value?) and that the sequencing of participatory reforms affects whether they are successful (on sequences, see Pierson 2004).

An important recent contribution by Gonzalez et al. (2020) reanalyzes Olken’s (2007) study of corruption in Indonesian road projects. They exploit the fact that different treatments in the original study were independently assigned and examine the interaction of top-down audits and bottom-up community monitoring in a multiplicative specification. In contrast to Olken (2007), they conclude that community participation did reduce missing funds, but only in the absence of simultaneous top-down audits, which dampen citizen participation in monitoring. They conclude that “the presence of effective and credible top-down monitoring can actually undermine the participation goals of a competing grassroots intervention” (Gonzalez et al. 2020: 27). Their estimates suggest that community monitoring can be as effective in reducing corruption as top-down audits, and almost three times as cost-effective.

Other studies provide less causal leverage. Two new studies from Brazilian municipalities offer insights about two outcomes that have received less attention in the literature. First, Touchton et al. (2020) compare municipal governments with participatory institutions (i.e., policy councils and participatory budgeting) to those without and find a positive correlation of participation with higher tax revenues. They mitigate endogeneity concerns by matching similar municipalities with and without such institutions. Interestingly, they find that this positive correlation between participatory budgeting and revenue collection appears to strengthen over time. Tracing the mechanisms of such reinforcing relationships might be one area of future research. In a second study from Brazil, Touchton and Wampler (2020) focus on the effect of participatory budgeting on a development outcome, infant mortality. They find the adoption of participatory budgeting is associated with lower infant mortality. Similar to their study on tax revenues, they also find that those municipalities who adopted participatory budgeting early had the largest effect sizes. While a more robust identification strategy is required to judge to what extent these are sorting effects, it nonetheless suggests that institutionalized participation might have positive general equilibrium effects in the long run that one-off interventions are less likely to achieve.

In a study from India, Costa-Font and Parmar (2021) find that being exposed to village health committees increases the use of maternal health services but not preventive child health care.

They mitigate endogeneity concerns by comparing births taking place before and after the reform. Importantly for our review, these committees develop health plans that have an impact on resource allocation. The study hints at spatial heterogeneity with larger villages and villages closer to district headquarters showing larger effect sizes. How such inequities can be addressed in low-capacity contexts might be subject to further research.

Further studies from the US (Calabrese et al 2020; Hagelskamp et al. 2020; Karner et al. 2019; Shybalkina and Bifulco 2019) consider how participation relates to the distribution of allocations across neighborhoods, projects, or budget categories. Several studies from South Korea (Grillos 2017; Hong and Cho 2018; Jung 2021; No and Hsueh 2020) investigate the impact of participation on allocations and aggregate outcomes. The study by Grillos (2017) considers participatory budgeting as a multi-stage process and finds that poverty-related bias stems from poorer neighborhoods being less likely to submit proposals. In contrast, Hong and Cho (2018) find that budget allocations for surveillance cameras were larger for low-income neighborhoods when the budget was decided through participatory rather than bureaucratic mechanisms. They also hint at a potential equity-effectiveness trade-off. Jung (2021) finds a positive association of participatory budgeting with measures of fiscal sustainability and administrative efficiency for municipal governments. Additionally, No and Hsueh (2020) highlight the importance of inclusive organizational structures for redistributive effects of participatory budgeting to materialize.

Finally, Saguin (2018) provides evidence from a community-driven development project in the Philippines. The study finds no impact on the poor's access to social services, their likelihood to participate in community activities, or their sense of solidarity. It does find a positive impact on per capita expenditure on the poor but points at the inability of the project to sustain participation by poorer citizens ultimately leading to elite control.

Overall, the review detects substantial growth in the literature on participation in budgeting. Among the themes that stand out, we note the reanalysis by Gonzalez et al. (2020), which fundamentally challenges one of Olken's (2007) central conclusions about the relative effectiveness of participation and transparency. Recent work also tackles important questions about the institutional and bureaucratic context into which reforms are introduced, the need to examine impacts in the long-run, and the sources of bias in participatory decision-making.

Revenues

Also welcome is new work linking open budgeting to revenue generation. Touchton et al. (2020), mentioned earlier, show that Brazilian municipalities that adopt participatory institutions collect more tax revenues. The study compares municipalities that created "policy councils" (where citizens can have a say on various aspects of local policy) and adopted participatory budgeting (devoting a share of the municipal budget to projects chosen by citizen assemblies) and municipalities that did not. Their estimates suggest municipalities collect 2.5 percent more tax revenues for each additional policy council they create, and an additional 16 percent when they adopt participatory budgeting, with that relationship becoming stronger over time. This suggests that the adoption of participatory institutions related to public policy and budgets can set off a virtuous cycle of good governance, with citizens paying more tax as they feel that they have a say in how governments spend public resources, and local governments collecting additional revenues to fund public services that citizens want.

The behavioral underpinnings of these findings are reflected in the results of an online survey experiment across 50 countries by Sjoberg et al. (2019). They find that “tax morale”, which refers to citizens’ willingness to voluntarily comply with taxes, increases when they are able to express their preferences on how governments should spend revenues, from 5.07 to 5.23 on a scale from 1 to 8 (p. 10). This study also stands out due to its geographic coverage. The results hold across this diverse sample including all continents.

While our focus has been on the effect of transparency or participation as independent variables, an inversed causal path is also worth considering. One area that illustrates this relates to revenues. Even though outside the scope of our review, Weigel (2020) provides evidence that a state’s attempt to tax citizens might induce participation. He found that in a Congolese city, in a treated neighborhood (i.e., where a door-to-door property tax collection campaign took place) not only did tax compliance go up, but citizens also increased their participation (by attending town hall meetings or submitting evaluations of government performance). This links to work reviewed above, pointing to a potential positive spiral of increased taxation capacity and participation even in contexts where state capacity is weak. Future research might explicitly focus on such reciprocal processes and trace their emergence with a focus on sequences and self-reinforcing processes (see Pierson 2004).

Procurement

Public procurement of goods, services, and public works accounted for US\$11 trillion of government spending across the world in 2018, equivalent to 12 percent of global GDP (Bosio and Djankov 2020). It is also an area that is frequently exploited for corruption (Rose-Ackerman and Palifka 2016). Hence, research looking at how open governments are with this specific type of expenditure, and what impacts more transparent procurement practices can have, is an especially welcome development in the recent fiscal openness literature.

Bauhr et al. (2019) examine procurement transparency in 30 European countries, drawing on more than 3.5 million government contracts awarded between 2006 and 2015, and look at how procurement transparency—both ex ante, during the call for tenders, and ex post, after the awarding of the contract—affects the risk of corruption. They measure the latter by noting whether contracts are awarded to a single bidder, even when markets are competitive. They show that more disclosure, especially during the call for tenders, improves horizontal accountability—monitoring by bidding firms—and reduces a contracting organization’s share of contracts awarded on the basis of a single bid, therefore reducing the likelihood of corrupt behavior. As single bidder contracts are on average more expensive, their numbers suggest that publishing more complete information during the call for tenders could result in potential savings of several billions of Euros annually.

As noted above, Zamboni and Litschig (2018) find that in Brazil, a randomized increase in the annual audit risk reduced corruption in procurement. The increase in audit risk they examine has a sizable impact although it is only about one-fifth of the increase examined in Olken’s (2007) field-defining study, suggesting that audits can contain corruption even with much lower audit probabilities. However, it is not clear to what extent this is attributable to the public release of audit information, as opposed to fear of detection and legal sanction, for instance. Finally, the study by Geradino et al. (2017), also mentioned above, cautions against unintended distortions in procurement where they induce organizations shift away from more transparent auctions towards less competitive direct contracting. This can in turn lead to a significant reduction in supplier competition and to a potential increase in prices of goods and services procured.

Beyond the award of tenders and procurement procedures, in many countries there are widespread concerns about poor services when delivery is outsourced. A pilot study by Keefer and Roseth (2021) of a program providing school meals in Colombia illustrates opportunities for research to examine this aspect. More work along these lines is required to assess the role of transparency and participation in ensuring outsourced services meet required standards.

4. Progress, stagnation, and next steps

Our review of recent research documents a vibrant and rapidly growing evidence base on the benefits of various elements of fiscal openness. In a short space of time, researchers have substantially increased the number of studies that can plausibly claim to identify causal effects, beyond the very small set of core results identified only a few years ago by de Renzio and Wehner (2017). The empirical context of research has also become much more geographically diverse, which is crucial for exploring the generalizability of findings in field-defining work by authors such as Ferraz and Finan (2008) and Olken (2007). Recent studies convincingly demonstrate that the power of audits travels well beyond the context of these initial studies. We also reported a growing number of studies of participatory budgeting that consider experiences beyond Brazil, where such practices were pioneered – although in this area most of the empirical work is correlational. In addition, we detected scholarly interest in previously neglected outcomes, notably revenues and procurement, including studies that identify causal effects (again from audits, on procurement). Overall, the study of the impacts of fiscal openness has become richer and more nuanced. The most well-documented causal effects are positive: lower corruption and enhanced accountability at the ballot box. The empirical case for open government in this policy area is rapidly growing in strength.

Nonetheless, there is plenty of scope to do more. Many aspects require substantial further investment to fully understand the impacts of fiscal openness and to design effective strategies to bring them about. While more causally identified evidence has been published or produced in a short space of time since the previous review, this evidence is very selectively targeted at a single, albeit important, component of fiscal transparency—ex post audits. In contrast, the pioneering work on Uganda by Reinikka and Svensson (2005, 2011) based on expenditure tracking surveys—linking improvements in the transparency of disbursements to local service delivery and education outcomes—appears not to have stimulated further publications in peer-reviewed journals. This is surprising, given that such surveys continue to be widely used by practitioners across various countries (e.g., Sundet 2007, 2008). More collaborations between practitioners and researchers in this area might generate new insights. A differentiated approach to examining authorized allocations and actual expenditures remains particularly important in contexts where gaps between planned and actual delivery of services or projects are often substantial, as is the case in many developing countries (e.g., Williams 2017).

With regard to the study of participation, the focus has been on participatory budgeting, where governments provide citizens with opportunities to influence allocations of projects. Yet far less attention has been paid to participation at different stages across the budgeting process, such as legislative approval or execution. The relationship between external audits and social audits—where citizens undertake an audit of service delivery in their locality—appears an important omission, too, as do other possibilities for interaction between external auditors and citizens (Mendiburu 2020). In many instances, it is not clear to what extent and how findings on participation in local processes might translate to regional or national policymaking. Moreover, we also noted uneven growth in academic attention to specific examples, in particular New York City. Overall, then, knowledge on the impacts of fiscal openness also remains partial and compartmentalized. We conclude with some reflections on future directions.

From local to national. Many of the most convincing research designs are deployed in decentralized settings, which facilitate quantitative analysis due to sample size and heterogeneity in local conditions and outcomes. Yet much funding, technical assistance, and advocacy focuses on the national level, which has crucial importance for public finances and development more broadly—especially in poor countries, which tend to be more centralized than industrialized democracies (Gadenne and Singhal 2014). There is thus a disjuncture between the localized focus of cutting-edge empirical research on the one hand, and the national-level focus of practitioners and donors on the other, where reform efforts have often been disappointing (Andrews 2013). However, quantitative evaluations of national reforms should also be possible. These might involve analyses of individual transactions instead of aggregate outcomes (e.g., Bauhr et al. 2019), or staggered rollouts of reforms that reach different institutional units at different times—Vannutelli (2021) adopts this strategy to study auditor assignment to municipalities, but similar approaches might be possible at the national level where reforms are rolled out across line ministries or other units in a staggered manner. Both disaggregation and staggered rollouts can help to overcome sample size constraints at the national level (by “splitting” the sample into smaller parts that vary in their treatment status either statically or across time), thereby allowing for more robust research designs. In addition, the impact of national-level events—such as the release of budget information—on public opinion, for instance in relation to trust or citizens’ evaluations of government performance, can be examined with surveys (as in Berliner and Wehner 2021). While there are challenges to documenting impacts of fiscal openness at the national level, there exist potential strategies to expand such research.

Working with governments. The past decade has seen substantial investment in field experiments on transparency and accountability, including with budget information. Many of these focus on information campaigns by civil society organizations or other researcher-led treatments, e.g. scorecards, flyers, SMS messages, etc. (e.g., Dunning et al. 2019). Some of this interest is driven by donors, while researchers might find civil society more willing to collaborate than governments (de Souza Leão and Eyal 2020). Yet it is unclear whether or to what extent evidence on information campaigns carried out by civil society or researchers allow any conclusions about the effectiveness or otherwise of government budget openness. To build the evidence base for the latter, we need more collaboration with governments to evaluate their actions directly. This might involve field experiments (e.g., Zamboni and Litschig 2018), or scholars working with governments in designing initiatives so that they facilitate learning about impacts, for example through staggered and randomized rollouts. Such collaborations can also support learning within government organizations and may feed directly into government decision making (Krause and Hernández Licona 2020).

From programs to systems. Budget systems consist of complex networks of institutions, such as finance ministries, line ministries, legislatures, and auditors, and across different levels of government. These institutions interact in structured ways and are embedded in a wider context (see also Mansoor and Williams 2021). Donors and public finance practitioners have invested in assessing the performance of budget systems, for example the Public Expenditure and Financial Accountability (PEFA) initiative. Yet causality is hard to pin down. Field experiments can help, but they adopt a program logic, and one-off interventions may not fully capture real-life processes. Natural experiments may better capture equilibrium effects, for instance where actions are at scale and information reaches citizens organically rather than in ways that are induced by researchers (Berliner and Wehner 2021). This may get closer to understanding accountability systems and how they evolve (IBP and INTOSAI Development Initiative 2020). For example, studies show that local media amplify the electoral accountability effects of audits (Ferraz and Finan 2008, Larreguy et al. 2020). More work is needed to integrate other components of budget systems, such as legislatures, and other institutions such as courts or anti-corruption agencies.