Preprint. Accepted for publication in IMPACT Journal.
The Version of Record is forthcoming and will be available from Sage Journals.

Commentary

The Unexamined Gate: Why the Search for Research Impact Keeps Looking in the Wrong Place

Christopher Chibwana[1]

Abstract

Research impact assessment misdiagnoses where impact fails. It focuses on the downstream work—treating impact as a delivery problem—when the decisive failure occurs upstream, in selecting which questions receive funding. This commentary argues that the funder has the most leverage over that choice and the least accountability for it. Competitive funding rewards proxies that run counter to what impact requires, compounding at the evaluation stage, where those same proxies punish the work impact demands. The vagueness of *impact* lets funders claim it as a mission while measuring it as outputs, with assessment regimes reproducing the blind spot, grading the research while exempting the choice that shaped it. Drawing on the author's vantage points and various evidence, the piece examines why even flexible, trust-based grantmaking does not escape this dynamic. Trust relocates the decisive choice rather than eliminating it, and the funder alone decides who to fund when proximate voices diverge. The commentary calls for shifting agenda-setting power to actors with lived experience and holding funders accountable for the legitimacy of that process, not individual grant outputs. The real test is not how generously a funder gives but how honestly it hands away the choice it most wants to keep invisible.

Keywords

Research impact, research agenda setting, funder accountability, research priority setting, research funding

Introduction

Years ago, as a researcher, I learned to write a fundable question. This is a craft, and like most crafts, it is mostly subtraction. You take the question that actually matters—the one that is tangled, slow to answer, and awkward for someone powerful—and you sand it down until it fits a 3-year grant, a clean identification strategy, and a reviewer's sense of what counts as rigor. I told myself I was translating the important into the researchable. Often, I was doing the opposite: letting the researchable stand in because the researchable is what got funded.

I then moved into the space between research and use. My job was to make knowledge usable: to dig out papers buried behind paywalls, distill them into policy briefs that a minister could read in 10 minutes, and wrap the findings in stories that made a room of decision makers lean in. I carried findings across the gap between those who produce knowledge and those who act on it.

I now sit on the other side of that table as a funder. From this seat, I can see the machinery I spent years responding to: the signals I read, the proxies I performed for, and the questions never written because no one would pay to answer them. I have become convinced the field has misdiagnosed where impact breaks down.

[1]William and Flora Hewlett Foundation, Menlo Park, CA, USA

Corresponding author:
Christopher Chibwana, William and Flora Hewlett Foundation, 2121 Sand Hill Road, Menlo Park, CA 94025, USA.
Email: cchibwana@hewlett.org

Preprint. Accepted for publication in IMPACT Journal.
The Version of Record is forthcoming and will be available from Sage Journals.

We spend enormous energy asking researchers to be more impactful, to disseminate, translate, and mobilize their knowledge. But impact rarely fails at the end of the pipeline. It fails at the beginning, at the question, and the funder is the actor with the most leverage over which questions get asked and the least scrutiny of its role.

Impact Fails Upstream, Not Downstream

Funders and the research-use community concentrate their effort at the far end of the process, after the research is done. We build translation units and knowledge-mobilization teams (the in-house functions that repackage findings for nonacademic users); fund policy briefs, stakeholder workshops, and dissemination plans; and ask how to get findings used. The implicit theory is that good research exists and the problem is delivery—that impact is lost in the *last mile* between a finding and the decision it should inform.

Some of it is lost there. But the more consequential loss happens before any research, in the choice of what to study. You cannot translate your way out of the wrong question. A beautifully disseminated answer to a question no one was asking is not deferred impact; it is foreclosed impact. The last mile gets attention because it is visible and measurable—a brief was published, a meeting held. The first mile, the selection of the question, is where the outcome is determined, and it is given almost no notice.

Medicine has measured its own waste. In a landmark audit, Chalmers and Glasziou (2009) estimated that around 85% of research investment is lost, and the first stage of loss, ahead of flawed design and unusable reporting, is the question itself: the wrong one, or one no patient or clinician was asking. The field's own reformers are explicit about who owns that stage: The primary responsibility for the losses that flow from decisions about what to research lies with the funder (Chalmers et al., 2014). The waste begins upstream and compounds downstream.

Funders, public and private, own that first mile. Researchers respond to the questions someone will pay to answer; the payoff structure is set upstream. Treating impact as a delivery problem, we examine the stretch of road we do not control and ignore the gate we do. The gate is the funder's, and it is where the leverage lives.

This is not naivety, nor for want of self-examination. A serious literature now scrutinizes the funder's role, much of it written by funders: Retrospective studies find academic excellence a poor predictor of wider impact (Guthrie et al., 2016), and funders have turned toward relational models and called openly for scrutiny of their own place in the system (Tseng et al., 2022). Much of this is admirable. But notice where it reaches—toward choosing better and getting research used. A funder can adopt all of it and still decide, alone, which questions are worth asking. The reflection has gone almost everywhere except there.

It Grades the Answer, Not the Question

Consider the most elaborate apparatus ever built to assess research. Each cycle, the Research Excellence Framework (REF) rates British universities' research to help allocate around £2 billion a year, judging outputs, environment, and societal impact (REF, n.d.). The impact component concerns us. Universities evidence it in case studies scored on reach and significance, and the verdict weighs heavily on careers and budgets.

The REF measures the wrong end of the process. It asks what difference a piece of research made, not whether it was the right research to fund. It grades the answer and exempts the question. The decisive choice—what was worth studying at all—was made years earlier, by the funder, in a room the REF does not

Preprint. Accepted for publication in IMPACT Journal.
The Version of Record is forthcoming and will be available from Sage Journals.

enter. The university is graded on impact, whereas the body that shaped the agenda is graded on nothing. The most influential actor in that chain faces the least scrutiny, and the measurement cannot see the gap: It can only assess research that exists. Questions ruled out before anyone could ask them leave no case study to score.

The Gate Selects Against Impact

Here is the uncomfortable part: The gate does not merely fail to select for impact; it often selects against it.

Watch what a competitive funding process rewards: the prestige of the applicant and their institution; citation potential and publication in the right venues; and a narrow, legible rigor—the clean design, the identifiable effect, the method a panel already recognizes. It rewards speed to a defined deliverable. Each is a proxy, easy to compare across applications and defend to a board.

Now consider what the work of impact looks like. It is slow, running on the timelines of institutions and relationships, not 3-year grant cycles. It is relational, built on trust that accrues over years and does not show up on output logs. It is messy and locally embedded, resistant to the clean design that scores well. Its results are hard to attribute because real change has many parents. Line the two lists up, and the conclusion is stark: The qualities that make a grant score well are close to the opposite of those that make research matter.

The failure then compounds at the other end. Suppose an impact-oriented grant slips through anyway. We evaluate it on outputs and milestones hit on schedule: the same proxies, now applied to judgment rather than selection. The grantee who spent 2 years building the relationship that will move a decision has little to show at the 18-month report. The lesson is learned quickly and rationally: Next time, produce the outputs, hit the milestones, and stop doing what impact requires.

So there are two distinct failures, worth naming separately because they have different fixes. There is a *selection failure*—choose the wrong questions, and people pass through the gate. And there is an *evaluation failure*—even the right work is judged by measures that punish it. A funder can be guilty of one, or both. Most are guilty of both.

The Vagueness Is Not Innocent

Why has this persisted? Partly because funders have never been made to say what they mean by impact.

Impact is the field's most-used and least-defined word, appearing in every strategy document with almost no operational definition. An undefined goal does not stay undefined; it collapses into whatever is easiest to count. Absent a shared meaning, impact becomes outputs: papers, briefs, citations, convenings. The measurable proxy fills the vacuum that the definition left behind.

This does more than mismeasure: It relocates the whole problem downstream. If impact just means the use of outputs, the interesting action is by construction at the point of use, exactly where the field has trained its gaze. The measurement manufactures the misdiagnosis.

The vagueness is convenient because it lets a funder hold impact as its stated mission while measuring outputs as its actual practice, without ever feeling the contradiction. The looseness of the word is not a harmless imprecision waiting for a better framework; it is what keeps the proxy game invisible to the public, to boards, and most of all to funders themselves.

Preprint. Accepted for publication in IMPACT Journal.
The Version of Record is forthcoming and will be available from Sage Journals.

The Honest Funder's Blind Spot

The alternative to selecting by proxy is to change how the choice is made: to fund flexibly and trust researchers to decide what is worth studying. This is right, and it removes a real harm; it lifts the output-and-timeline proxies off grantees, so the slow, relational work impact requires is no longer punished at evaluation. That is half the problem solved.

Flexible, trust-based funding relocates the decisive choice rather than eliminating it. A funder not selecting on citations or prestige is selecting on trust, and trust left unexamined is the oldest proxy of all. Fund the organizations you already know, and you have not escaped the gate; you have just stopped noticing you are the one holding it.

The difference is who the funder then chooses to trust. The approach worth defending prioritizes proximate actors: organizations led by people with lived experience of the problems their research addresses, trusted by decision makers because they share those problems rather than merely study them. Their access to power is not a separate credential to be gamed; it follows from their standing in the communities the research is about. Proximity to power is downstream of proximity to the problem. That is a proxy worth defending on exactly the grounds this essay cares about, selecting for the conditions of impact rather than against them.

And yet even this does not dissolve the funder's complicity. It concentrates it into a single, irreducible decision. A community is not a monolith. Among proximate actors, some will frame the problem differently, define the question differently, and pursue impact differently. Lived experience confers real standing to set the agenda, but it does not, on its own, tell the funder whose framing to fund when proximate voices diverge. That choice remains the funder's. It is the one part of agenda setting the funder cannot hand away and where, unchecked, all the old entrenchment reassembles under a better name.

A Fix That Constrains Funders, Not One That Frees Grantees

The instinct, having come this far, would be to prescribe flexible funding to everyone and be done. I want to resist that because it addresses the wrong side of the problem. Flexibility removes the evaluation squeeze, the output proxies imposed on grantees. It does nothing about selection and leaves the funder's own choices as unexamined as before. It frees the grantee while sparing the funder. Given that the funder is the under-scrutinized actor, that cannot be the answer. Necessary, but not sufficient.

The missing half must point scrutiny back at the funder, and it comes in two parts.

The first is to cede agenda setting, to build funding processes in which the questions originate with proximate actors rather than whoever the funder already trusts. This treats the selection failure at its root, moving the choice of what to study closer to the people it concerns. It directly addresses the entrenchment risk that even trust-based funding carries.

This is not a thought experiment. In British health research, the James Lind Alliance brings patients, carers, and clinicians together as equals to name the questions that matter to them, and across roughly a hundred conditions, its partnerships have repeatedly found that what patients most want answered is not what researchers choose to study (James Lind Alliance, 2021). But the same partnerships mark the limit of the move. People who live with the problem can set priorities that still go unanswered because nothing compels the research system to take them up (Staley et al., 2020). The question can be handed over even when the money that answers it has not been.

Preprint. Accepted for publication in IMPACT Journal.
The Version of Record is forthcoming and will be available from Sage Journals.

Philanthropy has built its own version. Participatory grantmaking cedes the decision itself—the strategy and the criteria, not just the short list—to the communities a funder means to serve, and a movement has gathered under the call to shift the power (Gibson, 2018; Global Fund for Community Foundations, n.d.). Some funds already run this way: The Red Umbrella Fund, directed by and for sex workers, decides through the people it exists to serve rather than for them (Red Umbrella Fund, n.d.). Even here the gate does not fully open: Staff still set who is eligible before the community chooses, so the last cut—whose framing counts—stays with the funder. The power can be shared; it has not yet been surrendered.

The first move has limits worth naming. Ceding agenda setting to proximate actors assumes a community that can be identified and convened. Where the affected population is dispersed, contested, or itself party to the problem, as in international peace and security, naming the proximate actor may be impossible without taking sides. There, ceding has little purchase, and almost everything rests on how openly and answerably the funder makes the choice itself.

The second closes the gap that ceding opens. If a funder simply hands the picks away, it can disclaim responsibility, which is another way of escaping scrutiny. So ceding the *what* must be paired with accountability for the *how*: The funder is answerable, across its whole portfolio, for the quality and legitimacy of the process by which it decides whom to trust and how it gives power away. Not "did your grants produce outputs," but "did you build a defensible process for deciding what to fund, and did you cede power honestly to a genuine range of proximate voices, rather than to a comfortable few?"

Together these resolve the tension raised earlier, that you cannot both cede the selection and be judged for it. The resolution is that the funder is no longer accountable for *which* grants are chosen but remains fully accountable for the design of the mechanism that chooses them—the decision that flexible funding, left alone, lets it make invisibly. Process accountability drags that back into the light: a scrutiny the funder cannot satisfy by pointing at a grantee's outputs, nor escape by pleading the choices were someone else's.

What would this ask of a funder, concretely? Four things, none of them unusual.

Set the agenda in the open. Before a call goes out, convene the people who live with the problem to name the questions worth funding—not to react to a list already drawn, but to draw it. The James Lind Alliance model shows this can be done with rigor.

Publish the basis of the choice. State the criteria by which questions and grantees are selected, and the reasoning behind them, so that agenda setting becomes a decision others can examine rather than one that happens by default.

Report the negative space. Account not only for what was funded but also for what was turned away and why—the questions declined, the framings not backed. A gate is known by what it stops, and that is exactly what funders never record.

Submit the process to outside review. A funder scrutinizes every grantee it backs; let its own selection be scrutinized in turn, by people with no stake in its continuation. Accountability a funder administers to itself is not accountability.

None of this asks a funder to fund worse research. It asks it to make the one choice it now makes privately—what is worth asking—answerably. Some funders can take all four steps; others only some, or only in part, depending on what they fund and on constraints that sit above any individual program officer. They are directions for travel, not a checklist. Partial progress counts, provided it moves toward sharing the choice rather than away from it.

Preprint. Accepted for publication in IMPACT Journal.
The Version of Record is forthcoming and will be available from Sage Journals.

Nor is any of this as simple as a funder letting go. A funder is rarely free: It answers to boards and governments, to the scientists it has long backed, and to the stakeholders expecting their issue to be funded this year. Each is a claim on the agenda, and ceding power to proximate actors means taking it from somewhere that will not always yield. Many hands hold the gate, not all of them the funder's. The point is narrower: Of those hands, the funder's hold the clearest leverage and require the least accountability. Systemic change will take more than funders, but the funder is where it can start, as the actor most able to act and least asked to.

Conclusion

I have argued this from the three seats I have occupied, and the view is consistent from all of them. As a researcher, I shaped my questions to the gate. As an intermediary, I watched good work die in the gap between what was fundable and what was needed. As a funder, I now hold the gate myself, and I can see how easily even a well-intentioned hand on it reproduces the failure it means to fix.

None of this makes funders villains. It makes them the actor the field has agreed not to look at. We have spent a decade improving everyone downstream of the money—the researchers, the translators, the users—while the one position with the most leverage over research impact has faced the least examination of its own role. That is no accident: It happens when the actor who could demand accountability is the one who would have to accept it.

So the test of a serious funder is not whether it funds flexibly, speaks the language of impact, or prioritizes the right kinds of grantees. It is whether it is willing to be held accountable for the decisions it most wants to keep invisible: what it chooses to fund, whom it chooses to trust, and how honestly it gives that power away. Everything else we measure is a way of looking somewhere easier.

Acknowledgements

During the preparation of this manuscript, the author used a generative AI assistant (Claude, developed by Anthropic) for drafting and editing support. The author conceived the argument, directed all substantive content, reviewed and revised all AI-assisted text, and takes full responsibility for the final manuscript. In accordance with the journal's policy on the use of AI, a disclosure generated with the sIFA tool accompanies this submission.

Author contributions

Christopher Chibwana: Conceptualization; Writing – original draft; Writing – review and editing.

Ethical considerations

This commentary did not involve human participants, human data, or human tissue; ethical approval was therefore not required.

Consent to participate

Not applicable.

Consent for publication

Not applicable.

Declaration of conflicting interest

The author is a Program Officer at the William and Flora Hewlett Foundation, a philanthropic grantmaking foundation. The commentary reflects the author's own views and experience as a grantmaker, not the positions of the Foundation, and concerns the role of funders in general rather than any specific grant or grantee. This position is disclosed as a potentially relevant interest.

Funding statement

The author received no specific grant for this work from any funding agency in the public, commercial, or not-for-profit sectors.

Data availability

Data sharing is not applicable to this article, as no new data were created or analysed.

References

Chalmers, I., Bracken, M. B., Djulbegovic, B., Garattini, S., Grant, J., Gülmezoglu, A. M., Howells, D. W., Ioannidis, J. P. A., & Oliver, S. (2014). How to increase value and reduce waste when research priorities are set. *The Lancet, 383*(9912), 156–165. https://doi.org/10.1016/S0140-6736(13)62229-1

Chalmers, I., & Glasziou, P. (2009). Avoidable waste in the production and reporting of research evidence. *The Lancet, 374*(9683), 86–89.

Gibson, C. M. (2018). *Deciding together: Shifting power and resources through participatory grantmaking* (J. Bokoff, Ed.). GrantCraft, Foundation Center.

Global Fund for Community Foundations. (n.d.). *#ShiftThePower*. Retrieved October 18, 2026, from https://globalfundcommunityfoundations.org/what-we-stand-for/shiftthepower/

Guthrie, S., Kirtley, A., Garrod, B., Pollitt, A., Grant, J., & Wooding, S. (2016). A 'DECISIVE' approach to research funding: Lessons from three Retrosight studies. *RAND Health Quarterly, 6*(1), 6.

James Lind Alliance. (2021). *The James Lind Alliance guidebook*. National Institute for Health and Care Research. https://www.jla.nihr.ac.uk/jla-guidebook/

Red Umbrella Fund. (n.d.).. Retrieved October 18, 2026, from https://www.redumbrellafund.org

Research Excellence Framework. (n.d.).. Retrieved October 18, 2026, from https://www.ref.ac.uk

Staley, K., Crowe, S., Crocker, J. C., Madden, M., & Greenhalgh, T. (2020). What happens after James Lind Alliance Priority Setting Partnerships? A qualitative study of contexts, processes and impacts. *Research Involvement and Engagement, 6*, Article 41. https://doi.org/10.1186/s40900-020-00210-9

Tseng, V., Bednarek, A., & Faccer, K. (2022). How can funders promote the use of research? Three converging views on relational research. *Humanities and Social Sciences Communications, 9*, Article 219. https://doi.org/10.1057/s41599-022-01157-w