Innovation Under Resource Constraints:

Supercomputing in Scientific Research

Justine Boudou* and John McKeon+

June 20, 2025

Abstract

Should organizations concentrate scarce resources among a few knowledge workers or distribute them more broadly? The answer to this question depends on whether additional resources enable more ambitious work or simply allow lower-priority projects to move forward. We examine how the direction of innovative output changes as resource constraints are relaxed in the context of high-performance computing (HPC), a critical input to modern scientific research. Using data from XSEDE, an NSF-funded program that allocates supercomputing resources to researchers, we leverage system-wide capacity constraints to identify the causal impact of variation in resource constraints on scientific output. We find that relaxing constraints increases the number of publications and shifts the direction of research. Scientists pursue less popular and newer topics, explore areas beyond their prior expertise, and broaden the scope of their work. However, these directional shifts are associated with fewer citations, suggesting a trade-off between frontier-expanding innovation and impact. Our findings show that allocation strategies shape not just the volume but the trajectory of innovation, with direct implications for R&D managers and policymakers supporting innovation under resource constraints.

Keywords: *Constraints*, *Innovation*, *Science*, *Resource Allocation*, *Research Direction*, *Knowledge Workers*

Authors contributed equally, listed alphabetically. *Boudou: Harvard Business School (jboudou@hbs.edu) +McKeon: Boston University (jmckeon@bu.edu). The authors thank Stephen Deems and Dave Hart for data access. We are indebted to Kevin Boudreau, Shane Greenstein, Josh Krieger, Megan MacGarvie, Raviv Muricano-Goroff, Kyle Myers, Felix Poege, Maria Roche, Tim Simcoe, Jeremy Watson, and Rosemarie Ziedonis for feedback on early drafts. We thank the participants of the BU Strategy & Innovation Brown Bag, TPRI seminar series, NBER Productivity Lunch, RISE Workshop, 2024 SMS Conference, MIT FutureTech and 2025 AIM Conference. All errors and omissions are our own.

1 Introduction

Many organizations, from public funding agencies and academic labs to firms managing R&D portfolios, must allocate scarce resources across competing projects (Boudreau et al. 2016, Blomfield & Vakili 2023, Mezzanotti & Simcoe 2023, Krieger et al. 2022, Nagaraj & Tranchero 2023). This leads to a central challenge: whether to spread these resources broadly among many researchers or concentrate them among a select few. This choice fundamentally depends on how researchers respond to changes in resource constraints – specifically, how they prioritize and select projects as resource constraints are relaxed. If additional resources enable researchers to pursue more exploratory or boundary-pushing projects, concentrating resources among fewer individuals may be preferable to foster breakthroughs. Conversely, if researchers merely use additional resources to expand along established trajectories or to pursue lower-priority incremental projects, distributing resources broadly might be more beneficial. Thus, this allocation decision hinges on a fundamental question: How does the direction of innovative output change as resource constraints are relaxed?

Existing literature offers mixed predictions about how resource constraints shape innovation (Ketkar & Roche 2024). Some argue that constraints sharpen focus and stimulate creativity (Mishina et al. 2004), while others suggest that relaxing constraints enables more exploratory or ambitious work (Cyert & March 1963, Hoegl et al. 2008). However, these perspectives often overlook the distinctive nature of knowledge work, which operates under different institutional logics and incentive structures than traditional production settings (Bikard 2018). Researchers typically enjoy substantial autonomy in setting their agendas and allocating effort, while also navigating complex professional incentives (Stephan 1996, Roach & Sauermann 2010, Stephan 2012, Gambardella et al. 2020). As a result, it is unclear how researchers respond when resource constraints are relaxed. On one hand, they may use additional resources to pursue ideas that were previously too risky or unfamiliar, helping them build a distinct scientific identity and differentiate their work

from prior scholarship (Dasgupta & David 1994, Foster et al. 2015, Uzzi et al. 2013). On the other hand, scientific recognition tends to reward depth and cumulative contributions within a specialized domain (Jones 2009), which may lead researchers to use additional resources to deepen existing lines of inquiry by producing follow-on, incremental work.

Empirically, isolating the causal impact of resource constraints on scientific production is challenging. Variation in resource constraints is rarely exogenous: it often coincides with changes in other inputs or is typically subject to selection, where more promising projects or researchers receive greater support. As a result, we lack clear empirical evidence on whether and how scientists change the nature of their work when constraints are relaxed. This paper addresses this issue by leveraging a novel empirical setting that allows us to study how variation in constraints on a critical input to innovation, supercomputing power, impact the innovative output of scientific researchers.

To do so, we assemble a novel dataset from the NSF-funded initiative known as *The Extreme Science and Engineering Discovery Environment* (henceforth "XSEDE"). XSEDE allows academic researchers to apply each quarter for access to a distributed network of high-performance computing resources across the United States. Despite the growing importance of high-performance computing (HPC) for frontier innovation across domains such as artificial intelligence, drug discovery, and climate modeling (Thompson et al. 2022, Kim 2025, Tranchero 2023), supercomputing capacity remains scarce and unevenly distributed (Lohr 2019), prompting increasing attention from both policymakers and firms. Over its ten-year span, XSEDE supported over $6 billion in funded research, enabling scientists to tackle computationally intensive scientific challenges that were previously beyond reach.[1] We combine internal XSEDE data on 710 scientific projects - including detailed records of computing allocations - with information on the publications linked to each project, allowing us to trace how variation in resource constraints shapes research output.

[1] See NSF Awards 1053575 and 1548562.

We exploit the allocation process of XSEDE to get plausibly exogenous variation in the degree to which researchers are constrained relative to the computing resources their projects require. In each cycle, researchers submit proposals requesting a specific amount of computing power. These requests are evaluated by a committee of domain experts, who assign a recommended allocation based on the amount of computing the project would require if resources were unconstrained - providing an objective benchmark of resource requirements. Throughout the period we study, demand for computing resources consistently exceeds supply. Hence, in each allocation cycle, the aggregate computing demand is then reconciled with available supply by reducing each recommendation based on an algebraic formula and allocating researchers this reduced amount. Our empirical framework leverages this unique resource allocation process of XSEDE in order to create variation in the quantity of computing resources received by researchers relative to what they would need, hence how constrained they are.

We find that increased resource allocation (i.e., a relaxing of constraints) leads to a substantial increase in the number of papers published within a project. Specifically, a 10 percentage point increase in allocation leads to an 11% increase in the number of papers written. More importantly, our results show that resource constraints also impact the direction of innovation. On average, we find that as resource constraints are relaxed, researchers tend to explore less studied (i.e., "unconventional") research avenues and broaden the scope of their research, pursuing projects that span a wider array of topics. These changes are driven respectively by the *least* conventional and the *broadest* papers within a researcher's output, indicating significant shifts at the boundaries of researchers' work. We also find that researchers explore new research streams among their *most* exploratory work, studying topics they have not previously studied. In addition, among the papers addressing the *newest* topics, researchers increasingly focus on topics that emerged more recently. These results suggest that the relaxing of constraints on resources enables researchers to push the

boundaries of their work, advancing into different territories. However, these shifts in research direction entail a significant trade-off. In particular, we also find that the relaxing of constraints leads to a decrease in the number of citations, though journal quality remains unchanged. This suggests that, while researchers use additional resources to expand the frontier of their work, their resulting contributions may be less useful to other scientists or slower to diffuse.

Taken together, our findings highlight a key strategic challenge for organizations which must allocate scarce innovative resources across several projects or knowledge workers. As resource constraints are relaxed, scientists tend to pursue more unconventional, exploratory, and broader projects - those they would likely deprioritize under tighter constraints. However, these changes carry a clear trade-off, as the work enabled by additional resources is less cited. These findings have important implications for allocation strategy. If the organizational goal is to maximize visibility or citations, distributing smaller amounts of resources across many researchers may be more effective, as highly-cited ideas tend to be pursued first. However, if the objective is to foster novel, more unconventional work that pushes the frontier, then allocating more substantial resources to fewer researchers may be necessary to enable that shift. More broadly, our findings underscore the importance of considering not only the rate but also the direction of scientific output when designing allocation strategies - whether in R&D-intensive firms, innovative startups, or publicly funded research programs that allocate scarce resources across projects.

This paper makes three contributions. First, it advances our understanding of how resource constraints shape innovative output among knowledge workers, where institutional logics and incentives differ from many other settings (Manso 2011, Azoulay et al. 2011). While prior work has linked variation in resource availability to changes in risk-taking, exploration, and creativity (Hoegl et al. 2008), it has largely abstracted away from the unique context of scientific knowledge production and the professional incentives individual scientists face. We bridge organizational theory and

science studies by embedding the constraint trade-off in this context (Stephan 2012, Jones 2009, Foster et al. 2015) and respond to a key limitation raised by prior work by focusing on directional changes in output (Agrawal et al. 2018). We show that as constraints are relaxed, researchers shift at the margin from more conventional, high-reward projects to those that are more exploratory or unconventional, revealing an organizational trade-off between impact and direction. These findings contrast with theories of slack resources fostering symmetric increases in both high- and low-impact output through a combination of effort and selection effects (Agrawal et al. 2018). Instead, we show that in the context of knowledge work, the organizational trade-off arises from a different margin of adjustment: when constraints are relaxed, scientists produce different kinds of work with lower visibility, leading to a trade-off between novel, more unconventional work and impact, rather than a dispersion in quality.

Second, we contribute to the literature on the production of scientific knowledge by providing new evidence on how researchers choices shape the nature of their output. Prior work has examined how individuals and firms increase risk-taking (Mandler 2017, Krieger et al. 2022, Franzoni & Stephan 2023), allocate effort across projects (Blomfield & Vakili 2023) and engage in search, recombination and collaboration (Fleming & Sorenson 2001, 2004, Fry et al. 2020, Fry 2023, Bikard et al. 2019, Vakili et al. 2022, Kaplan & Vakili 2015, Teodoridis 2018). Other studies explore how researchers depart from prior agendas (Myers 2020, Nagle & Teodoridis 2020, Hill et al. 2025, Lu 2024), expand existing work (Myers & Tham 2023), or respond to the introduction of new tools (Furman & Teodoridis 2020, Cavalli 2024), data infrastructures (Tranchero 2023, 2025), physical capital (Baruffaldi & Gaessler 2021) or different paces of change (Teodoridis et al. 2019). Our findings enrich this literature by showing how resource constraints influence the production of scientific knowledge - highlighting the role of resource allocation in shaping the scope, novelty, and trajectory of research. This complements prior work examining how tools and technologies shape

research direction by steering scientists toward particular questions (Tranchero 2023, Kim 2025, Teodoridis 2018, Furman & Teodoridis 2020); in contrast, we show how loosening constraints - without altering the tools - can expand the set of projects pursued and shift researchers' direction from within.

Finally, we provide causal evidence on how individual scientists respond to variation in resource constraints, using a novel empirical setting that offers a rare opportunity to observe plausibly exogenous variation within a real-world scientific environment. Empirical assessments of the relationships between resource constraints and innovation have been limited by their lack of a causal interpretation (e.g., Rosso 2014, Nohria & Gulati 1996, George 2005). We study the allocation of high-performance computing resources in the NSF-funded XSEDE program, which supports computationally intensive research across a wide range of scientific fields. By leveraging the program's allocation mechanism, we obtain plausibly exogenous variation in the amount of computing power researchers receive, relative to what their projects require. This allows us to estimate how resource constraints affect both the quantity and direction of scientific output.

2 Researcher Strategies Under Resource Constraints

Organizations engaged in R&D routinely face the challenge of allocating scarce resources across knowledge workers. As such, researchers often receive only a fraction of what they need to complete their research agendas. Let this constraint be formalized by $s \in (0, 1]$, denoting the share of resources (with respect to what they need) that researchers receive. In this section, we discuss researchers' marginal decisions as additional high-performance computing resources become available (i.e., as s increases): what kinds of research output do researchers undertake as resource constraints are relaxed?

2.1 Do researchers produce more output as resource constraints are relaxed?

We begin by examining how an increase in the amount of high-performance computing resources allocated to researchers (i.e., the relaxing of constraints) impacts the quantity of output they produce.[2] In practice, it is not obvious that more resources necessarily lead to an increase in output. Researchers could potentially be less efficient with additional resources, or they might leverage external alternatives, which could offset the expected gains. However, based on previous work and the limited availability of outside options, we might expect that an increase in the amount of computing resources leads to an increase in output.

First, existing literature tends to find a positive correlation between the availability of resources and scientific or innovation-related output, typically measured by publications or patents. For instance, access to data or to technology such as Information Technology (IT) is associated with an increase in the number of publications (Ding et al. 2010, Nagaraj & Tranchero 2023). Mezzanotti & Simcoe (2023) indicate that investments in basic research decline when funding costs increase. Additional work shows that tax credits stimulate additional R&D activity by reducing the cost of R&D (e.g., Hall & Van Reenen 2000). In the scientific context, research suggests that government funding leads to increased research output (Payne et al. 2003, Arora & Gambardella 2005, Ganguli 2017). These insights indicate that enhancing resource availability generally leads to greater research output.

Importantly, prior work suggests that the impact of getting more resources on research output largely depends on the availability of outside options. For instance, Jacob & Lefgren (2011) find that being funded by the National Institute of Health (NIH) has a positive but only small effect on the number of publications and citations and show that this is consistent with researchers being able

[2]We discuss the 'relaxing' of constraints to align with our empirical framework. Consistent arguments (in the opposite direction) could be made considering the 'imposing' of constraints.

to shift to other sources of funding. In contrast, high-performance computing resources typically lack direct substitutes, due to their specialized nature and the significant infrastructure required to support them which includes high costs and substantial space requirements (Dongarra et al. 2011, Millett & Fuller 2011). Indeed, unlike other research inputs that are fungible, HPC systems are uniquely designed to handle large-scale simulations, complex data analysis, and computationally intensive tasks. Further, large simulations often cannot be split up onto multiple resources but require an allocation within one environment. Hence, outside options are often limited. This suggests that an increase in the amount of computing resources will have a positive impact on researchers' output.

2.2 Do researchers produce a different type of output as resource constraints are relaxed?

A large body of research across disciplines has explored how resource constraints influence creativity and innovation (Acar et al. 2019). A central tension in this literature lies in whether constraints foster or hinder creativity and innovation-related outcomes. On one hand, some studies suggest that slack, i.e., a windfall of resources beyond immediate needs, enables organizations to take risk and engage in exploration by lowering the cost of failure (Cyert & March 1963, Hoegl et al. 2008). On the other hand, research - mostly in the context on financial constraints - finds that limited resources can stimulate innovation by focusing attention, reducing inertia, and encouraging more original problem-solving strategies (Mishina et al. 2004). The impact of constraints appears to be highly contingent on the type of constraint, the nature of the task, and the characteristics of the individuals or organizations involved (Rosso 2014, Acar et al. 2019). Yet, most of this literature centers on organizations (Katila & Shane 2005), team dynamics (Hoegl et al. 2008), or settings that differ significantly from professional knowledge work, often relying on empirical contexts with students or stylized tasks (Agrawal et al. 2018).

Understanding how knowledge workers - hereafter referred to as 'scientists' - respond to relaxed resource constraints requires taking into account the institutional logics that shape scientific decision-making (Bikard 2018, Bikard & Marx 2020). Scientists typically operate with considerable autonomy and limited managerial oversight, giving them substantial discretion in selecting and pursuing projects (Sauermann & Stephan 2013, Gambardella et al. 2020, Stephan 1996, 2012, Roach & Sauermann 2010). Given this autonomy, scientists' project choices are strongly guided by professional incentives embedded in the scientific institution, and also entail high opportunity costs (Azoulay et al. 2011, Manso 2011). Scientific recognition typically depends on producing publications and accruing citations, but also on establishing a distinct scientific identity and pursuing novel research that differentiates one's contributions from prior scholarship (Dasgupta & David 1994, Myers 2020). These features distinguish scientific knowledge production from the settings typically examined in prior work on resource constraints - such as firms, teams, students or laboratory experiments - where decisions are more managerial, tasks are more standardized, and incentives are more tightly coupled to short-term performance. As a result, how scientists respond to marginal changes in resource constraints may follow a different logic.

We examine how constraints on computing resources influence three key characteristics of the research output produced: (i) the conventionality of research topics (i.e., whether topics are widely studied or less common), (ii) the exploration of new research areas (i.e., whether researchers venture beyond their prior knowledge), and (iii) the breadth of individual papers (i.e., how many distinct topics they span). While these dimensions are not exhaustive, they capture important ways in which resource constraints might shape research direction. Conventionality speaks to how closely researchers remain to well-established lines of inquiry; exploration captures the degree to which researchers venture into unfamiliar intellectual territory; and breadth reflects how widely they draw from diverse areas of knowledge.

Conventionality of research topics - One possible effect of relaxing resource constraints is a shift toward less conventional research topics. When scientists face tight computing limits, they may prioritize safer projects with predictable outcomes, aiming to secure publishable results under constrained conditions. As constraints ease, the opportunity cost of pursuing riskier or more unorthodox topics declines. This may enable researchers to explore areas that are intellectually distinctive and less widely studied, albeit potentially more uncertain, slower-moving or more difficult to position for publication (Foster et al. 2015, Ketkar & Roche 2024). In this view, relaxing resource constraints grants scientists the flexibility to differentiate themselves through unconventional contributions that would otherwise be too costly or uncertain to pursue. Yet, this outcome is far from guaranteed. Scientists operate within institutional environments that reward conventional success, such as citation counts and alignment with established paradigms (Gius 2025, Hill et al. 2025). Unconventional research may carry real professional risks: smaller audiences, tougher peer review, and greater uncertainty about impact (Arts & Fleming 2018, Krieger et al. 2023, Piezunka & Dahlander 2015). Moreover, the public-good nature of scientific knowledge may further reduce incentives to explore less established areas: when knowledge is non-excludable, the returns from discovery of new research lines can be shared broadly, making it rational for individual scientists to free-ride on others' costly efforts to generate new knowledge (Hoelzemann et al. 2024). As such, even with additional resources, researchers may opt to pursue conventional topics that maximize career returns, leaving the net effect of relaxing constraints on research conventionality largely unchanged. Together, this leaves the impact of resource constraints on conventionality theoretically ambiguous.

Exploration - Constraints may also shape whether researchers venture beyond the bounds of their existing expertise or remain anchored in familiar domains. In theory, additional resources may allow researchers to test ideas or pursue topics outside their prior experience - exploration that may be

valuable but previously too risky, unfamiliar or costly (Levinthal & March 1981, Azoulay et al. 2011, Furman & Teodoridis 2020). However, exploration is not without costs in science: pivoting into a new area often entails a loss of accumulated expertise, a drop in productivity, and a reputational penalty for working outside one's core field (Hill et al. 2025, Arts & Fleming 2018). These costs are compounded by the structure of scientific careers, which tend to reward specialization, cumulative contributions, and the steady development of human capital within a narrow domain (Jones 2009). Thus, even with more resources, some researchers may avoid exploration, especially if doing so threatens their credibility, efficiency, or career momentum. Whether marginal increases in resources lead to more exploration is therefore ambiguous.

Breadth - Finally, constraints may affect the breadth of individual research projects - the extent to which it spans multiple distinct knowledge domains. Breadth, or scope, and its relationship with firm performance or innovation has been studied in various streams of literature, such as organizational learning (Dutt & Lawrence 2022, Ferguson & Carnabuci 2017). While prior research often uses knowledge breadth as a mechanism that impacts innovation (Nagle & Teodoridis 2020, Zhou & Li 2012), the scope of research may also be endogenous to the constraints researchers face. However, our understanding of how resource constraints impact the breadth of researchers' projects remains limited. On one hand, combining disparate knowledge sources is cognitively and logistically demanding, and such projects may simply be infeasible when resources are tight (Fleming 2001). In this case, relaxing constraints may allow researchers to initiate more ambitious, complex projects that require integrating a larger number of distinct ideas (Agrawal et al. 2018). Yet, the relationship between resource constraints and breadth is far from straightforward. Scientific institutions tend to reward deep specialization, with careers often built on sustained contributions to a single domain that aligns with disciplinary norms (Jones 2009). Broader projects may face practical and reputational barriers: they can be harder to publish in field-specific journals, more

difficult to review or evaluate, and may dilute the signal of expertise or novelty (Leahey et al. 2017, Bromham et al. 2016, Boudreau et al. 2016). Hence, the marginal resources researchers receive may simply be used to scale up depth - by refining analysis, increasing rigor, or extending existing lines of inquiry - rather than expanding scope. As a result, the relationship between resource constraints and breadth is theoretically ambiguous.

Overall, the relationship between resource constraints and research direction is theoretically unclear. While additional resources may unlock more unconventional, exploratory, or broader work, they may also reinforce more established, specialized projects aligned with researchers' existing expertise. We explore this tension empirically by examining how marginal increases in computing resource allocations (i.e., relaxing of constraints) affect the nature of the research scientists produce.

3 XSEDE

To investigate how constraints on supercomputing power impact innovation, we examine XSEDE, an NSF-funded initiative that provides a unique and valuable context for our study.

3.1 Background

XSEDE is an initiative funded by the National Science Foundation (NSF) in 2009 for a total of $257M.[3] XSEDE is one of the most advanced, powerful, and robust collections of integrated digital resources and services in the world for open scientific research and is designed to provide researchers across various disciplines with access to advanced computing resources. XSEDE establishes a distributed network of high-performance computing resources located across the United States. Figure 1 shows the location of the main supercomputers used as part of the program. Table A1 shows the ten most used XSEDE supercomputers during our study period. Their associated total capacity allocated during the period is expressed in the common metric used by XSEDE called

[3]XSEDE's predecessor was called TeraGrid and started in 2001. Following its success, XSEDE was replaced in 2022 by "ACCESS" (Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support).

Service Unit (SU), which can be roughly equated to "core-hours," and allows one to compare computing power across XSEDE resources whose computational specifications vary.[4]

Our interviews with several PIs highlighted how critical XSEDE is for their research agenda because it cannot be easily substituted by other options. One researcher we spoke with, who runs simulations to understand the impact of pollution and heat dispersion, emphasized that XSEDE is critical for their work because typical university clusters are too small. Another researcher, focused on using genetic material to unravel evolutionary history, noted that the computational demands of their work are so high that only XSEDE can provide the necessary resources, making it indispensable for their research. As such, XSEDE's computational resources have led to significant research across various disciplines by enabling complex and computationally intensive projects. For example, XSEDE has supported investigations into protein structures, where computational simulations facilitate the understanding of complex biological processes and disease mechanisms, aiding in the development of therapeutic strategies. Another example is in neuroscience, where researchers have utilized XSEDE to analyze large-scale single-cell RNA sequencing data to map neural circuitry and understand brain function, providing insights into cellular diversity and neural circuit mechanisms. Furthermore, researchers have used XSEDE resources to study how genetic diversity affects the survival rates of threatened animal species from youth to adulthood. These examples highlight XSEDE's critical role in advancing scientific understanding and innovation across an array of fields.

3.2 Resource Allocation Process

Research allocations are available to researchers affiliated with U.S.-based organizations, with the vast majority of these researchers being associated with universities or other research institutions.

[4]XSEDE converts machine specific core-hours to SUs using resource-specific conversion factors that are based on the LINPACK benchmark, a commonly accepted measure that captures a resource's ability to quickly solve problems.

Specifically, in our sample, 94.1% of researchers are associated with a university, 5.6% with a non-university research institution, and 0.3% with a commercial entity. Each quarter, the lead researcher - hereafter referred to as the Principal Investigator (PI) - can submit a proposal, which aims to tackle a specific scientific problem. These scientific problems are a bundle of related projects within a researcher's agenda. Proposals include a scientific background, research objectives, a computing allocation request, and a resource usage plan covering code performance timings, resource usage details, scaling information, and benchmark runs on the requested supercomputers (see an example of a proposal in the field of Biophysics in Appendix G).

For each of the quarterly windows of submission, the XSEDE Resource Allocation Committee (XRAC) is responsible for reviewing requests. On average, around 70% of projects are accepted and thus are awarded some computing power.[5] This committee includes about 60 computational researchers with expertise in various fields covering the diversity of fields of science encountered at XSEDE, with the most represented fields being more populated. Reviewers serve a three-year term. The role of this committee is to conduct an impartial and independent review of each resource allocation request in order to make a *recommendation* regarding the amount of resources that the project (or "grant") would need to be completed as proposed. Importantly for our empirical design, reviewers are asked to gauge requests only on their own merit (i.e., what does the project require?) and are specifically not informed of the overall quantity of resources available or the other requests across XSEDE's systems.[6]

In our study period, aggregate demand generally exceeds available supply. Hence, after each meeting, the "Allocations Team" meets to reconcile the recommendations made by the XRAC with the overall available supply of resources. This process starts with a "balancing" phase, where recom-

[5]Rejected projects are typically accepted in a future quarter, so there are very few outright rejections. This paper focuses on accepted proposals.

[6]Appendix H.1 provides excerpts from the *XSEDE Reviewer Manual* describing the review criteria for requests.

mendations from over-recommended resources are moved to similar ("architecturally equivalent") resources that are less over-recommended until the excess demand is balanced equally across similar computing platforms.[7] Next, this is followed by a "reconciliation" phase where recommendations are reduced through a formulaic solver in order to determine final *allocations*. The formula gives priority to NSF-funded projects, which represent about 40% of the proposals. Importantly, the only project-related characteristics considered in the reduction phase are (i) the recommendation made by the XRAC and (ii) the share of the project funding that is supported by the NSF. Otherwise, the reductions are a mechanical function of the level of over-subscription in each time period.[8]

Following an initial allocation, a PI can request additional resources in following quarter(s) via a "renewal." Renewal proposals are evaluated in the same way but also ensure that the prior allocation(s) led to some progress – we discuss this in more detail in section 4. Figure A1 summarizes the XSEDE allocation process. Figure 2 shows aggregate requests (i.e., demand), recommendations, and final allocations (i.e., reduced recommendations after considering supply available) for each quarter of our study period. Appendix H discusses more details on the allocation process.

4 Data

To study the impact of resource constraints on innovation within the XSEDE program, we use data about allocation awards covering each quarter between 2015 and 2022.[9] For each proposal, the data contains information at the supercomputer level about the initial request submitted by each PI, the recommended amount suggested by the XRAC, and the final allocation that is awarded. We drop resources corresponding to "storage" and we convert request, recommended, and allocated information associated with different supercomputers into standardized Service Units using conversion

[7] "Architecturally equivalent" refers to resources which have similar computational architectures and are thus near substitutes for the research that would be performed on them. For XSEDE, there are broadly two categories of architecturally similar resources: CPU-oriented systems and GPU-oriented systems.

[8] See Appendix H.2 for more details on the reconciliation formula.

[9] We are missing complete information about proposals for the years before 2015.

factors provided by XSEDE. We also observe information about the proposal field of science, as well as PI first name, last name, organization, and the percent of the funding coming from the NSF. Since both the initial request and any subsequent renewals are tied to the same project, we treat them collectively as a single XSEDE grant, associated with a PI. We therefore collapse information at the XSEDE grant (i.e., PI) level, keeping only grants for which we have complete information about their resources and for which we observe their first request. Each row of our dataset contains information about the total amount of requested resources made by the PI, the corresponding total amount recommended by the XRAC, and the total amount of allocated resources approved by the XSEDE allocation team. In theory, we could have kept our dataset at the request (i.e., quarter) level instead of collapsing it at the PI level. However, the scientific output that we observe can only be attributed to PIs and their XSEDE grants, and not to individual quarterly requests, making *grants* the relevant unit of observation. This corresponds to 1,062 unique grants.

In order to match each XSEDE grant to its scientific research production, we use two sources of data. We first perform a fuzzy match of our population of PIs to the *Dimensions AI* database using first name, last name and institution name.[10] Once PIs are identified, we retrieve their papers and keep those that reference XSEDE in the acknowledgments section.[11] Dropping book chapters, this gives us 3,659 papers. We also obtained access to the XSEDE internal publication database, where PIs can report the publications associated with their grant. From this, we retrieve 875 additional papers.

We then retrieve information about each paper using Dimensions AI. Our empirical analysis relies heavily on the notion of "concepts" from the Dimensions database. More precisely, the Dimensions

[10] See `https://www.dimensions.ai/`. In recent years Dimensions AI has been widely used in economics and management research, see e.g., Tian & Smith (2024), Myers et al. (2023), Myers & Tham (2023).

[11] From our discussions with PIs who were part of the XSEDE program, acknowledging XSEDE in the papers produced is the norm. We may be missing papers that attempt to acknowledge XSEDE without specifically mentioning a variation of "XSEDE" or "Extreme Science and Engineering Discovery Environment," but to be conservative we capture papers that we are confident originate from XSEDE computing allocations.

AI database uses machine learning techniques to derive concepts from papers' abstracts and rank them based on their relevance on a scale from 0 (not relevant) to 1 (very relevant).[12] In all analyses, we only consider concepts with a score above 0.5, in order to assign the most relevant topics to a given paper. Appendix I provides more detail on Dimensions concepts and our use of them.

Among the initial 1,062 grants, we observe scientific output for 719 of them. Among the 343 unmatched grants, 112 come from PIs we were unable to match or unable to uniquely match to Dimensions. For the remaining 231 grants, 1 grant has papers acknowledging XSEDE, but cannot be uniquely assigned to that grant because the associated PI is linked to multiple grants and 230 grants have PIs who did not report any paper to the internal XSEDE database nor acknowledge XSEDE in any of their publications.[13] Of the 719 grants with scientific output, we drop 9 grants whose ratio is strictly higher than 1 due to XSEDE data entry errors (see Table A2 for the successive sampling restrictions). Our final sample hence includes 710 unique grants. Section 5 discusses additional checks we perform on our data construction.

5 Empirical Strategy

Empirically, we want to estimate the impact of resource constraints on scientific output (rate and characteristics). In the ideal experiment, we would assign each researcher (i.e., grant) a random number of computing hours that is weakly less than what they require (i.e., randomly assign a level of constraint). We would then observe the subsequent output associated with the grant, which would allow us to estimate the effect of the marginal computing resource (i.e., marginal relaxing of constraint) on each researcher's output. While we cannot run such an experiment, we leverage the allocation process of XSEDE in order to get plausibly causal estimates.

[12] Several papers use MeSH terms to analyze the research topics of scientific papers (Azoulay et al. 2011, Myers 2020). MeSH terms are used for publications in PubMed which covers the biomedical literature. Conceptually, the use of Dimensions AI concepts allows us to extend this analysis to all the fields of science in our dataset.

[13] We drop these 230 grants from our data since we cannot identify which of these grants are "true zeroes" (i.e., did not produce any papers) versus those that did produce papers but did not cite XSEDE.

In most similar settings one could study, the typical threat to a causal interpretation is that projects and researchers with a higher expected productivity will receive higher computing allocations (will be less constrained), which could bias estimates. For instance, these projects would produce more papers regardless of the quantity of resources they receive, biasing our quantity estimate upwards. We overcome this by exploiting the fact that, conditional on the recommendation from the XRAC and the percentage of project funding that is supported by the NSF, each grant's final allocation is a function of the overall excess demand at the time of application, which is plausibly exogenous to the expected productivity of the grant.

Our regressions of interest are of the form:

$$Y_i = \beta_0 + \beta_1 \, \%Allocated_i + \delta \, \ln(Recommended_i) + \gamma \, \%NSF\,Funding_i + \mathbf{X}_i + \varepsilon_i \qquad (1)$$

where i indexes grants. Our treatment variable is $\%Allocated_i = \frac{Allocated_i}{Recommended_i}$ which captures the fraction of resources allocated to a grant relative to what it would have objectively needed as assessed by the XRAC – higher values indicate a grant is less constrained. This specification allows us to interpret the coefficient β_1 as the impact of receiving more computing resources relative to what a researcher would have needed.[14] Note that by definition of the XSEDE Allocation process, this ratio is less than or equal to one (i.e., all projects are weakly constrained). We control for the recommendation made by the XRAC with the logarithm of $Recommended_i$. One can think of this variable as capturing a combination of project quality and scale, where larger projects will inherently exhibit different output quantity and characteristics than smaller projects. The variable $\%NSF\,Funding_i$ accounts for the share of the project funding that is supported by the NSF, since this is part of the reconciliation formula. The vector $\mathbf{X}_i$ represents a set of fixed effects.

[14] In Appendix E, we confirm all results are robust to using the level of allocation, $\ln(Allocated_i)$, as the treatment variable. Appendix F shows results where we allow for $\%Allocated_i$ to have a non-linear effect.

We include field of science fixed effects, since some fields might apply for specific resources which are deferentially more or less busy and these fields might also have different research norms (e.g., more papers or different research characteristics).[15] We also include grant start year fixed effects in order to control for changes in overall supply over time as well as the potential for changing research output and characteristics over time. Our model includes the interaction of field of science and grant start year fixed effects which control for field specific time trends. We use robust standard errors clustered at the grant start year $\times$ field level because the reconciliation process is performed across architecturally similar machines.[16]

We perform several tests to assess whether our specification provides us with plausibly causal estimates. First, we confirm that observable characteristics of researchers do not correlate with the level of constraint a project has. In Table A3, we show that the main independent variable of interest, $\%Allocated$, is uncorrelated with pre-determined covariates related to the PI: seniority, institution, number of papers published before XSEDE, and number of citations received before XSEDE.

Second, given that the XRAC has information about the amount of computing resources requested by researchers, one might worry that recommendations are "anchored" on researchers' requests, leading some researchers to receive higher recommendations because they requested more resources than what their projects needed. In order for this to create endogeneity issues, asking for more should be correlated with both (i) the outcomes (i.e., PIs who ask for more should be either more or less productive than the average, or produce research with different characteristics) and (ii) the main explanatory variable (i.e., PI asking for more should have a higher or lower $\%Allocated$ than the average). While we cannot observe which PIs strategically asked for more than their

[15] Our data has three different levels of field of science aggregation: directorate (least granular), division, and field (most granular). Our baseline results use field but our results are robust to using the other levels of aggregation.

[16] Fields of science at a given point in time often require particular types of computing resources that are architecturally similar, leading to a reconciliation process that may be correlated within field-year.

true need, we create a proxy for the share of over-requested resources by measuring the "percent over-request" ($\frac{requested_i - recommended_i}{recommended_i}$), since larger gaps between the recommendation and request *may* suggest that the request has been inflated by the researcher. We then examine the relationship between this variable and $\%Allocated$, including the controls. Table A4 shows no significant correlation. To further assess this concern, we reached out to the XSEDE leadership team for their qualitative assessment. They emphasized that the system is designed to be robust against strategic manipulation. Proposal reviewers are domain experts who are aware of such possibility and actively account for them in evaluation. Review criteria require detailed, methodologically grounded justifications and emphasize computational efficiency - all within a strict page limit. As the XSEDE team noted, "it's very difficult to slide through bonus work, inefficient work, or unnecessarily expanded work." Moreover, large requests - the most obvious candidates for inflation - are more heavily scrutinized and typically face greater reductions during the reconciliation process (see Appendix H.2). As one member put it, "the game isn't nearly as easy to play or win as it might sound." Overall, while we cannot rule out all strategic behavior, the structure of the review process and the multiple frictions it imposes make it unlikely that such behavior occurs systematically or in a way that would bias our estimates.[17]

6 Results

6.1 Descriptive Statistics

Table 1 shows descriptive statistics about our sample of 710 grants. On average, researchers request 61.6 million service units (SU). The XRAC makes recommendations equal to 41.7 million SU on average and researchers gets a final average allocation of 35.8 million SU, with large variation in

[17] We also revisit checks on our data construction in light of our empirical framework. Note that a PI's decision to report to the internal XSEDE database is voluntary. We confirm that there is no correlation between our treatment variable ($\%Allocated$) and papers identified through the use of the internal database. Additionally, in order to assess whether grants in our final population are a selected sample, we construct an indicator equal to 1 if the grant is part of the 710 grants that constitute our main sample and 0 otherwise and we regress this measure on the main independent variable, $\%Allocated$, and the main controls. Table A5 shows no statistically significant relationship.

this number. On average, the share of allocated resources compared to the recommendation made by the review committee is 90%, but this variable is skewed. This can be seen in Figure 3 which shows the histogram of the treatment variable (*% Allocated*). On average, each grant is associated with 6.3 papers. Figure A2 shows the 15 most common fields of science in our sample. The most represented field is Materials Research with 56 grants, followed by Biophysics (49 grants) and Materials Engineering (42 grants).

6.2 Quantity of Science

In order to study the impact of additional resources on *quantity*, our outcome of interest is the log of the total number of publications associated with each grant.[18] Table 2 shows the results, where we regress this outcome on the share of allocated resources received by grants. Column (1) shows a naïve regression including only %*Allocated* and year fixed effects and displays a strong positive relationship, significant at the 5% level. Columns (2) and (3) successively add controls for the amount of computing resources recommended by the XRAC and the share of funding coming from the NSF, both of which enter the reconciliation formula. Column (4) also controls for field of science fixed effects and Column (5) adds grant start year $\times$ field fixed effects. Using Column (5) as our preferred specification, results show a strong positive relationship significant at the 1% level: a 10 percentage point increase in the percentage of allocated resources leads to approximately an 10.9% increase in the number of papers produced. Equivalently, an increase from the 10^{th} to the 75^{th} percentile in the share of allocated resources leads to a 52.9% increase in the number of papers produced.[19] Column (6) confirms that results are robust to using a Poisson specification with the count of papers as the outcome. This result implies that scientific output increases as resource

[18] Since our outcome does not contain zeroes (each grant in our sample is associated with at least one paper), we avoid issues described in Chen & Roth (2024).

[19] The 10th percentile of %*Allocated* equals 59% and the 75th percentile of %*Allocated* equals 1.

constraints are relaxed.[20]

6.3 Type of Science

In what follows, we explore the consequences of an increase in the share of allocated computing resources on the type of science that researchers pursue. We successively study whether such increase leads to: (i) researchers pursuing less traveled research avenues (ii) researchers studying topics beyond their expertise and (iii) researchers broadening the scope of their research.

6.3.1 Conventionality

First, we study whether an increase in the share of computing resources impacts researchers' propensity to work on less-studied topics. For each paper, we calculate the number of times each concept it entails has been used by other PIs in the same field, excluding all papers from the focal PI, and average this measure at the paper level.[21] The higher this number, the more the paper entails concepts that have been used by other researchers too – these are "more studied" concepts. The lower this number, the more the paper entails concepts that have been less studied by other researchers – these are "less frequently-studied" concepts. We then collapse this to the grant level using different measures of interest.[22] First, we look at the average among papers within grant, which gives us information about how well-studied the topics of the grant are on average. We also gain further insights about the overall distribution by considering the minimum among papers within grant, i.e., the average use of concepts by other researchers for the least "well-studied" (most "unconventional") paper. Finally, we also consider the maximum among papers within grant, i.e., the average use of concepts by other researchers for the most "well-studied" paper.

[20] One may be concerned about "strategic relabeling" whereby researchers migrate work to XSEDE that would otherwise be completed using outside resources. If this were the case, we would expect to observe a decrease in non-XSEDE papers by the researcher during the XSEDE time period. In Table A6, we confirm that this does not appear to be a systematic issue.

[21] We normalize this number to account for publication year following Perry & Reny (2016).

[22] We use the log of this measure. In Table A7, we instead create indicator variables equal to one for grants above median of this measure and zero otherwise which allows us to keep observations whose value is zero. We find consistent results.

Table 3 shows the results, with Column (1) using the average as the outcome, Column (2) using the minimum as the outcome and Column (3) using the maximum as the outcome. All columns control for the amount of resources recommended by the XRAC, as well as the share of funding coming from the NSF. Column (1) displays a negative and statistically significant coefficient, showing that the higher the share of allocated resources, the less popular are the topics studied on average. The coefficient implies that a 10 percentage point increase in the share of allocated resources leads to the study of concepts that are about 6.3% less studied by other researchers. Such a relationship could be driven by changes at the tails (e.g., the papers that study the least (or most) popular topics explore more niche topics) which would inform us about the boundaries of researchers' work. Column (2) shows that this result is driven by papers in the left tail of the distribution: for the paper with the least studied topics, a 10 percentage point increase in allocation leads to the study of concepts that are about 16.3% less studied by other researchers. Column (3) shows no change at the right tail of the distribution. Overall, we find that an increase in the share of computing resources leads researchers to explore less-studied concepts, driven particularly by shifts in papers that focus on the least-studied topics.

Thus far, results consider topic usage both before and after a given grant's start date. However, this negative relationship could be driven by (i) the study of topics that were initially less popular but rising in popularity (ii) the study of topics that were popular but losing attractiveness or (iii) the study of topics that have a persistently lower level of use by others. In order to investigate this, we begin by calculating for each paper the number of times each concept it entails has been used by other PIs *prior to*, and separately *after*, the paper publication year and average this measure at the paper level. This allows us to study a relative trend in use by others before versus after publication year to disentangle (i)-(iii). We then reiterate our previous analysis, considering as outcome the mean value of this measure across all papers associated with a grant, the minimum

value of this measure across all papers associated with a grant, and the maximum value of this measure across all papers associated with a grant. Results are presented in Columns (1) to (3) of Table A8 using the '*use prior*' measure and show similar results. Results presented in Columns (4) to (6) of Table A8 use the '*use after*' measure and again show similar findings. This suggests that when researchers are allocated a higher share of resources, they study topics that are persistently less popular – i.e., less conventional.

We next study whether the concepts used are more or less recently introduced to the scientific literature as researchers get allocated additional computing resources – e.g., are these "less studied" concepts also "newer to the world?" We take each concept associated with a paper and calculate the first year we see this concept appearing among XSEDE researchers.[23] We then average this measure at the paper level. This gives us information about the average 'vintage' of the concepts used in a paper. The higher this measure, the more recent the concepts it includes are. The lower this measure, the older the concepts it studies are. Column (4) of Table 3 uses as an outcome the average of this value across all papers associated with a grant. This measure proxies how recent the concepts used by a grant are on average. Columns (5) and (6) use as the outcome respectively the minimum and maximum of this value across all papers associated with a grant, which gives us information about the paper with the oldest and the most recent concepts, respectively. Column (4) shows no significant effect, implying that an increase in the share of computing resources does not impact the *average* age of concepts associated with a grant. We also do not find any significant result in Column (5), implying that an increase in the share of computing resources does not impact the age of concepts associated with the "oldest vintage" papers. However, Column (6) shows a positive and statistically significant estimate at the 5% level, implying that an increase

[23]Since we do not observe the *universe* of scientific articles in our data, we cannot observe with certainty when the topic was initially studied overall. We use the first time we observe a topic in XSEDE as a good proxy for its appearance in the computational scientific literature. Further, we have no reason to suspect a correlation between when concepts appear among XSEDE researchers relative to the universe and a grant's computing allocation.

in the share of computing resources leads to the use of more recent concepts associated with the "newest vintage" papers. This suggests that an increase in the share of computing resources leads researchers to advance their work into newer territories, as it is not the "oldest vintage" papers that change, but rather the "newest vintage" ones that explore even newer ideas. More specifically, Column (6) shows that a 10 percentage point increase in $\%Allocated$ leads to the study of topics that are about 0.47 years newer on average.

Taken together, we find strong evidence consistent with the fact that as resource constraints are relaxed, researchers choose to study topics that are less common and newer to the world.

6.3.2 Exploration

An increase in the share of computing resources could also impact researchers' propensity to "pivot," i.e., to work on topics that they haven't studied previously. To analyze this, we calculate for each paper in our sample the share of concepts that have not been used by the focal PI in the years preceding their entry into XSEDE. We then study the average, minimum, and maximum of this measure at the grant level. Conceptually, we are capturing the degree of "pivoting" or "exploration" by each researcher, studying topics that they are less familiar with. Using the average allows us to capture the general tendency of researchers to explore new topics across all papers within the grant as computing resources increase. Using the minimum within a grant helps us see what happens for papers that study the most familiar topics. Using the maximum gives us information about papers that focus on the most unfamiliar topics, i.e., the most "exploratory" papers.

We use the pool of papers written by the researcher in the two years preceding their entry into XSEDE to define the initial pool of papers with which we compare papers written as part of XSEDE.[24] Table 4 shows the results. Column (1) examines the mean share within grant. Column (2) examines the minimum and Column (3) examines the maximum. We find no significant effects

[24]Table A9 reiterates this analysis using a 3-year and 5-year pre-XSEDE period and finds similar results.

on the mean nor minimum measures but strong positive results on the maximum value within grant. More precisely, a 10 percentage point increase in the share of allocated resources leads to about a 1.2 percentage point increase in the share of new concepts. Overall, our findings imply that as resource constraints are relaxed, researchers explore 'new-to-them' topics among the most exploratory of their papers. This finding is consistent with the idea that the relaxing of computing constraints allows researchers to pursue more "exploratory" work that they do not seem to pursue in more constrained environments.

6.3.3 Breadth

Finally, an increase in the share of computing resources could impact the breadth of research projects pursued by researchers. To explore this, we study the effect of an increase in the share of allocated computing resources on the number of concepts used by papers associated with a grant. We first calculate the number of unique concepts associated with each paper, which informs us about the diversity of ideas it embodies. Appendix I shows an example of this measure. Our first outcome averages this measure at the grant level. We use this as a proxy for the breadth of research associated with a grant: the higher the average number of concepts, the broader the research on average. We also use as an outcome the minimum of this value across all papers associated with a grant. We use this as a proxy for the breadth of the "narrowest" paper associated with a grant. Similarly, we use as an outcome the maximum of this value across all papers associated with a grant. We use this as a proxy for the breadth of the "broadest" paper associated with a grant.[25]

Table 5 shows the results. Column (1) studies the average within grant and shows a significantly positive coefficient, implying that a 10 percentage point increase in the share of allocated resources leads to a 4.5% increase in the number of concepts associated with a grant. This indicates that

[25]In contrast to the journal breadth measure used in Azoulay et al. (2011), our measure more directly captures researchers' pursuit of broad topics separately from the impact in the scientific community.

a higher share of allocated resources leads researchers to expand their research breadth. Column (2) studies the minimum within grant and shows no effect. Column (3) uses as an outcome the maximum within grant and shows that the result in Column (1) is driven by the "broadest" papers becoming broader. This implies that as resource constraints are relaxed, researchers work on broader papers and specifically expand the broadest ones.

6.4 Journal Quality and Citations

Thus far, we find evidence that as resource constraints are relaxed, researchers write more papers and change the direction of their research. These findings provide insight into how computing resources influence researchers' strategies, as they reflect deliberate choices undertaken by researchers. In this section, we explore downstream outcomes that the researchers have less control over. We use journal placement as our measure of *ex-ante* quality and normalized forward citations as our measure of scientific impact.

First, we study the impact of additional resources on the quality of journal placement for papers associated with a grant. Specifically, we exploit information about the SNIP (Source Normalized Impact per Paper) value of the journal associated with each paper. SNIP measures the average citation impact of the publications in a given journal, which we use as a proxy of journal quality independent of the citations accrued to any given paper.[26] Column (1) of Table 6 examines the average SNIP value of papers associated with a grant. Columns (2) and (3) repeat this exercise by considering respectively the minimum and maximum SNIP value of papers associated with a grant. Overall, we do not find any significant estimate, suggesting that an increase in computing resources may not have a significant impact on the *ex-ante* quality distribution. This suggests that an increase in the share of allocated resources does not lead to a change in inherent quality.

Next, we study the impact of an increase in the share of computing resources allocated to re-

[26]SNIP values are normalized to account for differences in citation norms over years.

searchers on citations. We use an adjusted citation measure that normalizes by publication year and the field of science associated with each paper.[27] Since this measure already accounts for field differences, we do not control for field fixed effects or field × year fixed effects in these analyses.[28] This essentially allows us to compare the citations received by papers in our sample relative to the *universe* of papers, while normalizing by publication year and field. We then collapse this measure at the grant level, calculating the mean, minimum, and maximum across papers. Results are presented in Table 6. Column (4) uses as the outcome the average number of citations received by papers associated with each grant. Columns (5) and (6) repeat this exercise by using the number of citations associated with the least (respectively most) cited paper associated with a grant. Columns (4) and (6) show non-statistically significant estimates, while Column (5) shows a strong negative relationship. This suggests that, while an increase in the share of computing resources does not lead to a significant change in average scientific impact or reception, it does lead to a significant decrease in the number of citations received by the least-cited papers associated with a grant.[29] Given the results in columns (1)-(3) and our findings about changes in research direction, we are cautious to interpret this finding as representative of lower quality papers *per se*. While this is possible, this finding may be a consequence of the changing characteristics of the work produced which may correlate with less follow-on work and thus citations by other researchers.

To further explore this finding, we shed light on the association between the directional changes we estimate and this citation decrease. Table A10 regresses normalized citations on our different measures capturing research direction. Column (1) shows that an increase in the use of conventional

[27]This measure is provided by Dimensions AI and calculates the scientific influence of a paper by dividing the number of citations it received by the average number of citations received by papers published in the same year and field of research. A value of one means that a paper receives exactly the average number of citations expected for its publication year and field.

[28]We still control for grant start year FE since the Dimensions normalization is by paper publication year, which is different than grant start year.

[29]In unreported analyses, we construct an alternative field-year normalized citation measure that only uses our sample rather than the universe of papers and find similar results.

topics is associated with an increase in citations. Column (2) shows that an increase in the use of more recent topics is associated with a decrease in citations. Column (3) shows that an increase in the share of new concepts is associated with a decrease in citations. Given that our previous results concluded that an increase in the share of allocated resources leads to the study of less conventional, more recent, and more exploratory topics (see Figure 4 for a summary of the results), Columns (1), (2) and (3) are consistent with the idea that these changes in research direction lead to less citations. Column (4) shows that an increase in breadth is associated with an increase in citations. Given that our previous results concluded that an increase in the share of allocated resources leads to broader research, this finding suggests that broader research does not contribute to the observed decline in citations.

7 Discussion

Organizations engaged in innovation routinely face decisions about whether to concentrate limited resources among a few researchers or distribute them broadly. The trade-offs involved fundamentally depend on how individual scientists respond to changes in resource constraints, particularly whether additional resources enable them to pursue more exploratory, unconventional projects, or reinforce established lines of research or lower-quality work. A large body of research has examined how constraints influence creativity, risk-taking and organizational performance (Acar et al. 2019, Rosso 2014). Yet, relatively little is known about how such constraints influence the decisions of individual scientists, who operate with substantial autonomy and face distinctive professional incentives that simultaneously reward specialization and distinctiveness. As a result, when resource constraints are relaxed, it remains unclear whether researchers will shift direction or delve deeper into their specialized domain.

In this paper, we address this gap by examining how variation in constraints on supercomputing resources influences the nature of scientific output. Our empirical setting is XSEDE, a national

program that allocates supercomputing resources across research projects. Because demand for these resources consistently exceeds supply, XSEDE employs a centralized allocation process that reduces researchers' requested resources based on system-wide capacity constraints. By leveraging this process, we observe plausibly exogenous variation in the share of computing resources granted relative to what researchers would need, enabling us to estimate how constraints affect both the quantity and direction of scientific output.

We find that relaxing resource constraints leads to a significant increase in the number of papers published. More importantly, we observe systematic shifts in the direction of research. Scientists who are less constrained pursue less popular and more recently emerging topics, venture beyond their prior areas of expertise, and broaden the conceptual scope of their work. This suggests that under tighter constraints, researchers focus on safer, more conventional projects, and expand the boundaries of their work when constraints are relaxed. Importantly, these directional shifts are particularly salient for the *least* conventional papers, those that study the *newest* topics, the *most* exploratory papers, and the *broadest* papers, respectively – implying the relaxing of these constraints actually drives researchers to "spread their wings" and push the boundaries of their work. However, these shifts come with trade-offs. In particular, relaxing resource constraints leads to a decrease in citations, despite no change in journal placement.

These findings have important implications for managers and policymakers responsible for deciding whether to concentrate scarce resources among a few individuals or distribute more broadly. Our results highlight that such decisions affect not only the quantity of research produced, but also its direction - shaping which ideas get pursued and which remain sidelined. When researchers are constrained, they tend to prioritize more conventional and high-returns projects; when constraints are relaxed, they are more likely to pursue work that is more exploratory, unconventional, or integrative, albeit with lower returns. This suggests that allocation strategies should be guided

by organizational goals: if the objective is to support more boundary-pushing, field-expanding research, allocating more substantial resources to fewer individuals may be necessary. By contrast, if the goal is to sustain productivity along established lines, distributing smaller allocations more broadly may be more effective. Strategically timing the relaxation or tightening of constraints may allow organizations to better manage the trade-off between novelty and productivity. For example, relaxing constraints at early stages may enable researchers to explore more freely, while imposing constraints later may help focus effort (Aghion et al. 2008).

At a system level, our findings raise questions about whether current resource allocation practices overemphasize efficiency at the expense of exploration. If resource constraints consistently steer researchers toward safer, more conventional work, some potentially transformative ideas may remain perpetually underfunded and underexplored. This may have consequences not only for knowledge advancement, but also for downstream innovation. Firms often draw on academic research when developing new technologies, and the kinds of ideas scientists choose to pursue can influence which technologies are ultimately developed and commercialized (Marx & Fuegi 2020, Arora et al. 2018, Aghion et al. 2008, Cohen et al. 2002, Bikard & Marx 2020). As such, resource allocation decisions within the scientific community have ripple effects across the broader innovation system, shaping the knowledge base that fuels future technological progress.

Notably, the shifts in research output we highlight occur in the absence of any explicit changes to formal incentives. Researchers operate within the same professional and institutional environments, yet the nature of the work they pursue changes as constraints are relaxed. This complements prior work such as Azoulay et al. (2011), which shows that shifts in research direction can also result from changes in incentive structures. Taken together, these studies suggest that both resource availability and incentive design play distinct and potentially complementary roles in shaping the trajectory of scientific research.

This paper contributes to a deeper understanding of how constraints shape scientific production by offering new insight into how individual scientists prioritize projects when resources are limited. While previous research has shown that constraints can influence organizational creativity, we bring this conversation into the domain of scientific knowledge production, where individuals operate with substantial autonomy and face unique professional incentives. Conceptually, we link literatures on resource constraints and scientific incentives, highlighting a previously underexplored trade-off: relaxing constraints enables researchers to pursue more novel, unconventional, or integrative directions, but these shifts come at a cost in impact. Empirically, our setting offers a rare opportunity to study how individual scientists respond to resource constraints using plausibly exogenous variation in the allocation of a critical input - computing power - within a real-world scientific environment. By leveraging the allocation mechanism of a national supercomputing program, we are able to isolate the effect of marginal changes in resource availability on both the quantity and direction of research output. This allows us to provide causal evidence on how resource constraints shape not only how much science gets done, but what kind of science researchers choose to pursue.

Our findings suggest several promising directions for future work. While we focus on computing resources, future research could examine whether similar dynamics arise with other critical inputs and how multiple constraints interact to shape scientific behavior. Finally, while we study short- to medium-term outcomes such as publications and citations, future work could examine the longer-run effects of constraint relaxation, including delayed impact, translational outcomes, or contributions to downstream innovation. This could help further clarify the trade-offs organizations face when allocating scarce resources for innovation.

References

Acar, O. A., Tarakci, M. & Van Knippenberg, D. (2019), 'Creativity and innovation under constraints: A cross-disciplinary integrative review', *Journal of Management* **45**(1), 96–121.

Aghion, P., Dewatripont, M. & Stein, J. C. (2008), 'Academic freedom, private-sector focus, and the process of innovation', *The RAND Journal of Economics* **39**(3), 617–635.

Agrawal, A., Catalini, C., Goldfarb, A. & Luo, H. (2018), 'Slack time and innovation', *Organization Science* **29**(6), 1056–1073.

Arora, A., Belenzon, S. & Patacconi, A. (2018), 'The decline of science in corporate R&D', *Strategic Management Journal* **39**(1), 3–32.

Arora, A. & Gambardella, A. (2005), 'The impact of NSF support for basic research in economics', *Annales d'Economie et de Statistique* pp. 91–117.

Arts, S. & Fleming, L. (2018), 'Paradise of novelty—or loss of human capital? exploring new fields and inventive output', *Organization Science* **29**(6), 1074–1092.

Azoulay, P., Graff Zivin, J. S. & Manso, G. (2011), 'Incentives and creativity: evidence from the academic life sciences', *The RAND Journal of Economics* **42**(3), 527–554.

Baruffaldi, S. & Gaessler, F. (2021), 'The returns to physical capital in knowledge production: Evidence from lab disasters', *Max Planck Institute for Innovation & Competition Research Paper* (21-19).

Bikard, M. (2018), 'Made in academia: The effect of institutional origin on inventors' attention to science', *Organization Science* **29**(5), 818–836.

Bikard, M. & Marx, M. (2020), 'Bridging academia and industry: How geographic hubs connect university science and corporate technology', *Management Science* **66**(8), 3425–3443.

Bikard, M., Vakili, K. & Teodoridis, F. (2019), 'When collaboration bridges institutions: The impact of university–industry collaboration on academic productivity', *Organization Science* **30**(2), 426–445.

Blomfield, M. & Vakili, K. (2023), 'Incentivizing effort allocation through resource allocation: Evidence from scientists' response to changes in funding policy', *Organization Science* **34**(1), 100–128.

Boudreau, K. J., Guinan, E. C., Lakhani, K. R. & Riedl, C. (2016), 'Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science', *Management Science* **62**(10), 2765–2783.

Bromham, L., Dinnage, R. & Hua, X. (2016), 'Interdisciplinary research has consistently lower funding success', *Nature* **534**(7609), 684–687.

Cavalli, G. (2024), 'How academic labs adapt to advances in AI: The impact of alphafold', *Working Paper* .

Chen, J. & Roth, J. (2024), 'Logs with zeros? Some problems and solutions', *The Quarterly Journal of Economics* **139**(2), 891–936.

Cohen, W. M., Nelson, R. R. & Walsh, J. P. (2002), 'Links and impacts: the influence of public research on industrial R&D', *Management Science* **48**(1), 1–23.

Cyert, R. M. & March, J. G. (1963), *A Behavioral Theory of the Firm*, Blackwell, Cambridge, MA.

Dasgupta, P. & David, P. A. (1994), 'Toward a new economics of science', *Research Policy* **23**(5), 487–521.

Ding, W. W., Levin, S. G., Stephan, P. E. & Winkler, A. E. (2010), 'The impact of information technology on academic scientists' productivity and collaboration patterns', *Management Science* **56**(9), 1439–1461.

Dongarra, J., Beckman, P., Moore, T., Aerts, P., Aloisio, G., Andre, J.-C., Barkai, D., Berthou, J.-Y., Boku, T., Braunschweig, B. et al. (2011), 'The international exascale software project roadmap', *The International Journal of High Performance Computing Applications* **25**(1), 3–60.

Dutt, N. & Lawrence, M. (2022), 'Learning to manage breadth: experience as repetition and adaptation', *Organization Science* **33**(4), 1300–1318.

Ferguson, J.-P. & Carnabuci, G. (2017), 'Risky recombinations: Institutional gatekeeping in the innovation process', *Organization Science* **28**(1), 133–151.

Fleming, L. (2001), 'Recombinant uncertainty in technological search', *Management Science* **47**(1), 117–132.

Fleming, L. & Sorenson, O. (2001), 'Technology as a complex adaptive system: Evidence from patent data', *Research Policy* **30**(7), 1019–1039.

Fleming, L. & Sorenson, O. (2004), 'Science as a map in technological search', *Strategic Management Journal* **25**(8-9), 909–928.

Foster, J. G., Rzhetsky, A. & Evans, J. A. (2015), 'Tradition and innovation in scientists' research strategies', *American Sociological Review* **80**(5), 875–908.

Franzoni, C. & Stephan, P. (2023), 'Uncertainty and risk-taking in science: Meaning, measurement and management in peer review of research proposals', *Research Policy* **52**(3), 104706.

Fry, C. V. (2023), 'Crisis and the trajectory of science: Evidence from the 2014 ebola outbreak', *Review of Economics and Statistics* **105**(4), 1028–1038.

Fry, C. V., Cai, X., Zhang, Y. & Wagner, C. S. (2020), 'Consolidation in a crisis: Patterns of international collaboration in early covid-19 research', *PloS one* **15**(7), e0236307.

Furman, J. L. & Teodoridis, F. (2020), 'Automation, research technology, and researchers' trajectories: Evidence from computer science and electrical engineering', *Organization Science* **31**(2), 330–354.

Gambardella, A., Khashabi, P. & Panico, C. (2020), 'Managing autonomy in industrial research and development: A project-level investigation', *Organization Science* **31**(1), 165–181.

Ganguli, I. (2017), 'Saving soviet science: The impact of grants when government R&D funding disappears', *American Economic Journal: Applied Economics* **9**(2), 165–201.

George, G. (2005), 'Slack resources and the performance of privately held firms', *Academy of management Journal* **48**(4), 661–676.

Gius, L. (2025), 'Commercializing contrarian ideas: Evidence from AI contests', *Working Paper* .

Hall, B. & Van Reenen, J. (2000), 'How effective are fiscal incentives for R&D? A review of the evidence', *Research Policy* **29**(4-5), 449–469.

Hill, R., Yin, Y., Stein, C., Wang, D. & Jones, B. (2025), 'Adaptability and the pivot penalty in science', *Nature* .

Hoegl, M., Gibbert, M. & Mazursky, D. (2008), 'Financial constraints in innovation projects: When is less more?', *Research Policy* **37**(8), 1382–1391.

Hoelzemann, J., Manso, G., Nagaraj, A. & Tranchero, M. (2024), 'The streetlight effect in data-driven exploration', *National Bureau of Economic Research Working Paper* .

Jacob, B. A. & Lefgren, L. (2011), 'The impact of research grant funding on scientific productivity', *Journal of Public Economics* **95**(9-10), 1168–1177.

Jones, B. F. (2009), 'The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder?', *Review of Economic Studies* **76**, 283–317.

Kaplan, S. & Vakili, K. (2015), 'The double-edged sword of recombination in breakthrough innovation', *Strategic Management Journal* **36**(10), 1435–1457.

Katila, R. & Shane, S. (2005), 'When does lack of resources make new firms innovative?', *Academy of Management Journal* **48**(5), 814–829.

Ketkar, H. & Roche, M. (2024), 'Too Much, Too Soon? Early Funding, Technological Unconventionality, and Innovation Capabilities'.

Kim, S. (2025), 'Navigating the Rugged Data Landscape: The Impact of Data-Extrapolation Technologies on Knowledge Production', *Working Paper* .

Krieger, J. L., Myers, K. R. & Stern, A. D. (2023), 'How important is editorial gatekeeping? Evidence from top biomedical journals', *Review of Economics and Statistics* pp. 1–33.

Krieger, J., Li, D. & Papanikolaou, D. (2022), 'Missing novelty in drug development', *The Review of Financial Studies* **35**(2), 636–679.

Leahey, E., Beckman, C. M. & Stanko, T. L. (2017), 'Prominent but less productive: The impact of interdisciplinarity on scientists' research', *Administrative Science Quarterly* **62**(1), 105–139.

Levinthal, D. & March, J. G. (1981), 'A model of adaptive organizational search', *Journal of economic behavior & organization* **2**(4), 307–333.

Lohr, S. (2019), 'At tech's leading edge, worry about a concentration of power', *The New York Times* .

Lu, J. (2024), 'Intellectual distance and propensity to engage with new technological development: Evidence from electric vehicle technologies', *Working Paper* .

Mandler, M. (2017), 'The benefits of risky science', *The Economic Journal* **127**(603), 1495–1526.

Manso, G. (2011), 'Motivating innovation', *The Journal of Finance* **66**(5), 1823–1860.

Marx, M. & Fuegi, A. (2020), 'Reliance on science: Worldwide front-page patent citations to scientific

articles', *Strategic Management Journal* **41**(9), 1572–1594.

Mezzanotti, F. & Simcoe, T. (2023), 'Research and/or development? Financial frictions and innovation investment', *Working Paper* .

Millett, L. I. & Fuller, S. H. (2011), *The future of computing performance: game over or next level?*, National Academies Press.

Mishina, Y., Pollock, T. G. & Porac, J. F. (2004), 'Are more resources always better for growth? resource stickiness in market and product expansion', *Strategic Management Journal* **25**(12), 1179–1197.

Myers, K. (2020), 'The Elasticity of Science', *American Economic Journal: Applied Economics* **12**(4), 103–134.

Myers, K. R., Tham, W. Y., Thursby, J., Thursby, M., Cohodes, N., Lakhani, K., Mural, R. & Xu, Y. (2023), 'New facts and data about professors and their research', *Working Paper* .

Myers, K. & Tham, W. Y. (2023), 'Money, time, and grant design', *Working Paper* .

Nagaraj, A. & Tranchero, M. (2023), 'How does data access shape science? Evidence from the impact of US Census's research data centers on economics research', *Working Paper* .

Nagle, F. & Teodoridis, F. (2020), 'Jack of all trades and master of knowledge: The role of diversification in new distant knowledge integration', *Strategic Management Journal* **41**(1), 55–85.

Nohria, N. & Gulati, R. (1996), 'Is slack good or bad for innovation?', *Academy of Management Journal* **39**(5), 1245–1264.

Payne, A. A., Siow, A. et al. (2003), 'Does federal research funding increase university research output?', *Advances in Economic Analysis & Policy* **3**(1), 1–24.

Perry, M. & Reny, P. J. (2016), 'How to count citations if you must', *American Economic Review* **106**(9), 2722–2741.

Piezunka, H. & Dahlander, L. (2015), 'Distant search, narrow attention: How crowding alters organizations' filtering of suggestions in crowdsourcing', *Academy of Management Journal* **58**(3), 856–880.

Roach, M. & Sauermann, H. (2010), 'A taste for science? phd scientists' academic orientation and self-selection into research careers in industry', *Research Policy* **39**(3), 422–434.

Rosso, B. D. (2014), 'Creativity and constraints: Exploring the role of constraints in the creative processes of research and development teams', *Organization Studies* **35**(4), 551–585.

Sauermann, H. & Stephan, P. (2013), 'Conflicting logics? a multidimensional view of industrial and academic science', *Organization science* **24**(3), 889–909.

Stephan, P. (2012), *How economics shapes science*, Harvard University Press.

Stephan, P. E. (1996), 'The economics of science', *Journal of Economic Literature* **34**(3), 1199–1235.

Teodoridis, F. (2018), 'Understanding team knowledge production: The interrelated roles of technology and expertise', *Management Science* **64**, 3625–3648.

Teodoridis, F., Bikard, M. & Vakili, K. (2019), 'Creativity at the knowledge frontier: The impact of specialization in fast- and slow-paced domains*', *Administrative Science Quarterly* **64**, 894–927.

Thompson, N. C., Ge, S. & Manso, G. F. (2022), 'The importance of (exponentially more) computing power', *Working Paper* .

Tian, T. Y. & Smith, E. B. (2024), 'Stretched thin: How a misalignment between allocation and valuation underlies the paradox of diversity achievement in higher education', *Administrative Science Quarterly* **69**(3), 711–746.

Tranchero, M. (2023), 'Finding Diamonds in the Rough: Data-Driven Opportunities and Pharmaceutical Innovation', *Working Paper* .

Tranchero, M. (2025), 'Data-driven search and the birth of theory: Evidence from genome-wide association studies', *Working Paper* .

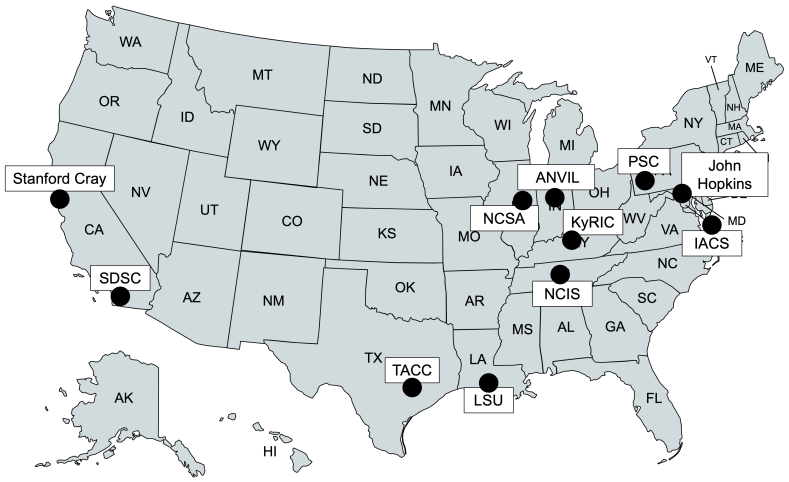
Uzzi, B., Mukherjee, S., Stringer, M. & Jones, B. (2013), 'Atypical combinations and scientific impact', *Science* **342**(6157), 468–472.

Vakili, K., Teodoridis, F. & Bikard, M. (2022), 'Detrimental collaborations in creative work: Evidence from economics', *Organization Science* **33**(5), 1741–1755.

Zhou, K. Z. & Li, C. B. (2012), 'How knowledge affects radical innovation: Knowledge base, market knowledge acquisition, and internal knowledge sharing', *Strategic Management Journal* **33**(9), 1090–1102.

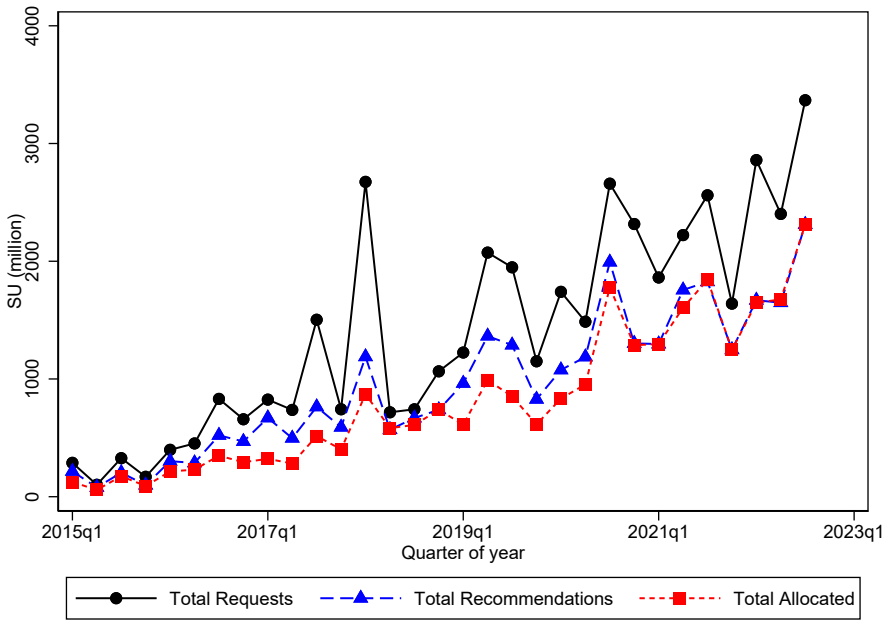
Figures

Figure 1: Location of XSEDE resources



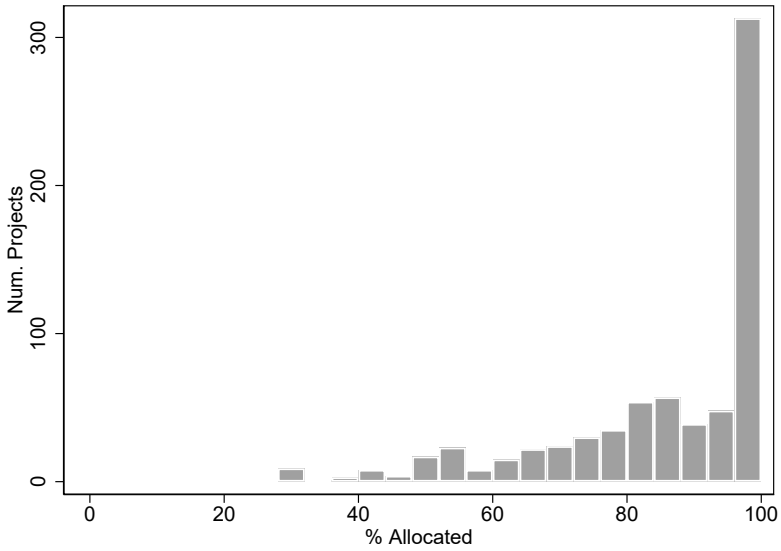
Notes: This figure shows the location of the main supercomputers that are associated with XSEDE.

Figure 2: Total Request, Recommendation and Supply Over Time



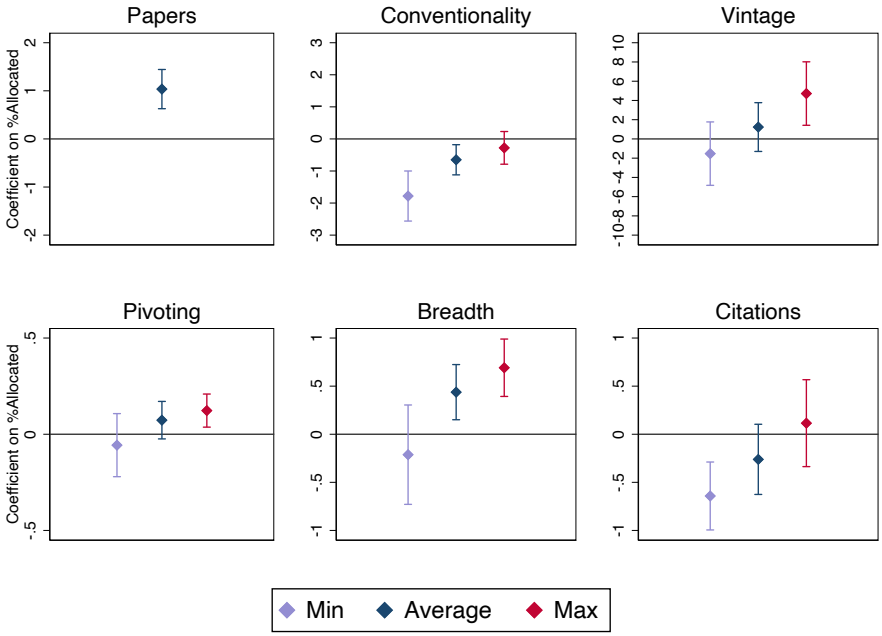
Notes: The x-axis shows each quarter between 2015 and 2022. 'Total Requests' represents the total amount of requested resources, 'Total Recommendations' represents the total amount of recommendations made by the XRAC committee, and 'Total Allocated' represents the total amount of allocated resources after reconciliation of supply and demand. Series show aggregate demand across all XSEDE machines. Several resources (Indiana Jetstream, Kentucky Research Informatics, PSC Bridges-2, Purdue Anvil) were launched in 2021, leaving the XSEDE supply less constrained and explaining the near-overlap between the 'Total Recommendations' and the 'Total Allocated' in the latest quarters.

Figure 3: Histogram of the share of resources allocated



Notes: This figure reports the distribution of the treatment variable *%Allocated*, which represents the ratio of computing resources allocated to a grant compared to the amount of resources recommended by the XRAC.

Figure 4: Coefficient Estimates - Summary



Notes: Each point represents the coefficient estimate on *%Allocated*. Regressions are run using the specification of Equation 1. Each dependent variable used is first calculated at the paper level and is described in Section 6. 'Average' takes the average among papers within grant. 'Min' takes the minimum among papers within grants. 'Max' takes the maximum among papers within grants.

Tables

Table 1: Summary statistics

	Mean	SD	Min	Max
Requested (Mil. SU)	61.6	144.3	0.0	1627.6
Recommended (Mil. SU)	41.7	91.2	0.1	930.7
Approved (Mil. SU)	35.8	80.2	0.1	862.9
% Allocated	0.9	0.2	0.3	1.0
Grant start year	2017.9	2.0	2015.0	2022.0
%NSF Funding	0.4	0.5	0.0	1.0
R1 universities	0.8	0.4	0.0	1.0
Junior PI	0.2	0.4	0.0	1.0
Number of papers	6.3	9.2	1.0	97.0
log(use)	1.0	0.8	0.0	11.8
Mean concept year	2004.4	4.7	1980.0	2023.0
Share of new concepts (2yr)	0.8	0.2	0.0	1.0
Average number of unique concepts	24.8	9.7	1.0	66.7
Observations	710			

Notes: All summary statistics are at the grant level. *Log (use)*: average number of times concepts used in papers associated with a grant are used by other researchers, see details in section 6.3.1. *Mean concept year*: average year in which concepts used in papers associated with a grant are first observed, see details in section 6.3.1. *Share of new concepts (2yr)*: constructed by calculating for each paper the share of concepts that have not been used by the focal PI in the two years preceding their first entry into XSEDE, see details in section 6.3.2. *Average number of unique concepts*: average number of unique concepts used in papers associated with a grant, see details in section 6.3.3.

Table 2: Quantity of papers

	DV = ln(papers)					DV = Num. Papers
	(1)	(2)	(3)	(4)	(5)	(6)
$\%Allocated_i$	0.532**	0.962***	0.940***	0.725***	1.036***	0.789**
	(0.209)	(0.199)	(0.208)	(0.217)	(0.246)	(0.365)
$\ln(Recommended_i)$		0.277***	0.275***	0.283***	0.285***	0.315***
		(0.022)	(0.023)	(0.027)	(0.036)	(0.044)
%NSF $Funding_i$			0.038	0.030	-0.072	-0.043
			(0.075)	(0.085)	(0.104)	(0.152)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field FE				Yes	Yes	Yes
Field × Year FE					Yes	Yes
Specification	OLS	OLS	OLS	OLS	OLS	Poisson
Dep Var Mean	1.29	1.29	1.29	1.31	1.35	6.58
R-Sq	0.05	0.24	0.24	0.36	0.46	
Observations	710	710	710	710	710	710

Notes: The unit of observation is a grant. In Columns (1)-(5), the outcome is the log number of papers associated with a grant. Column (6) shows a Poisson specification where the outcome is the number of papers associated with a grant. $\%Allocated$ is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. $\ln(Recommended)$ is the log of the number of SUs recommended by the XRAC. $\%NSF\ Funding$ is the share of funding coming from the NSF. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Less traveled research avenues

	DV = log(Use)			DV = Concept Vintage		
	(1)	(2)	(3)	(4)	(5)	(6)
Paper w/i grant:	Avg. Paper	Least Conv.	Most Conv.	Avg. Paper	Oldest	Newest
$\%Allocated_i$	-0.650**	-1.781***	-0.279	1.234	-1.529	4.717**
	(0.284)	(0.470)	(0.308)	(1.531)	(1.988)	(1.991)
$\ln(Recommended_i)$	0.054	-0.134**	0.154***	-0.360**	-1.278***	0.443**
	(0.048)	(0.059)	(0.051)	(0.179)	(0.244)	(0.219)
%NSF $Funding_i$	-0.022	-0.087	-0.050	-0.615	-0.626	-1.016
	(0.118)	(0.146)	(0.135)	(0.470)	(0.680)	(0.675)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Dep Var Mean	-0.41	-1.27	0.07	2003.85	2000.22	2007.54
R-Sq	0.67	0.55	0.63	0.44	0.44	0.39
Observations	613	529	613	665	665	665

Notes: The unit of observation is a grant. In Columns (1) to (3), we first calculate for each paper the average number of times the concepts it entails have been used by other PIs in the same field. Column (1) takes the mean of this value at the grant level, logged. Column (2) takes the min of this value at the grant level, logged. Column (3) takes the max of this value at the grant level, logged. In Columns (4) to (6), we calculate for each concept associated with a paper the first year we see it appearing and take the average at the paper level. Column (4) averages this measure at the grant level. Column (5) takes the min of this measure at the grant level. Column (6) takes the max of this measure at the grant level. $\%Allocated$ is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. $\ln(Recommended)$ is the log of the number of SUs recommended by the XRAC. $\%NSF\ Funding$ is the share of funding coming from the NSF. Standard errors are clustered at the grant start year × field level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Topics beyond expertise

	DV = Share of new concepts (vs prev. 2y)		
	(1)	(2)	(3)
Paper w/i grant:	Avg. Paper	Lowest Share New	Highest Share New
$\%Allocated_i$	0.073	-0.057	0.123**
	(0.059)	(0.099)	(0.052)
$\ln(Recommended_i)$	-0.002	-0.043***	0.014**
	(0.008)	(0.015)	(0.007)
%NSF $Funding_i$	0.006	0.021	-0.006
	(0.026)	(0.044)	(0.023)
Year FE	Yes	Yes	Yes
Field FE	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes
Dep Var Mean	0.78	0.63	0.88
R-Sq	0.45	0.37	0.42
Observations	600	600	600

Notes: The unit of observation is a grant. For each paper, we calculate the share of 'new-to-the-researcher' concepts, i.e., concepts that have not been previously used by the PI in the 2 years that precede XSEDE entry. Column (1) takes the average of this measure at the grant level. Column (2) takes the min of this measure at the grant level. Column (3) takes the max of this measure at the grant level. $\%Allocated$ is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. $\ln(Recommended)$ is the log of the number of SUs recommended by the XRAC. $\%NSF\ Funding$ is the share of funding coming from the NSF. Standard errors are clustered at the grant start year × field level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Research Breadth

	DV = log(Unique Concepts)		
	(1)	(2)	(3)
Paper w/i grant:	Avg. Paper	Narrowest Paper	Broadest Paper
$\%Allocated_i$	0.438**	-0.213	0.691***
	(0.173)	(0.312)	(0.180)
$\ln(Recommended_i)$	-0.009	-0.258***	0.047*
	(0.023)	(0.047)	(0.024)
%NSF $Funding_i$	0.037	0.150	0.015
	(0.046)	(0.125)	(0.053)
Year FE	Yes	Yes	Yes
Field FE	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes
Dep Var Mean	3.09	2.32	3.37
R-Sq	0.35	0.41	0.36
Observations	653	653	653

Notes: The unit of observation is a grant. We calculate for each paper the number of unique concepts it entails. Column (1) averages this measure at the grant level, logged. Column (2) takes the min of this measure at the grant level, logged. Column (3) takes the max of this measure at the grant level, logged. $\%Allocated$ is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. $\ln(Recommended)$ is the log of the number of SUs recommended by the XRAC. $\%NSF\ Funding$ is the share of funding coming from the NSF. Standard errors are clustered at the grant start year × field level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Journal quality and number of normalized citations

	DV = log(SNIP)			DV = log(Norm. Cites)		
	(1)	(2)	(3)	(4)	(5)	(6)
Paper w/i grant:	Avg. Paper	Lowest SNIP	Highest SNIP	Avg. Paper	Least Cited	Most Cited
$\%Allocated_i$	-0.111	-0.256	-0.019	-0.261	-0.642***	0.115
	(0.154)	(0.176)	(0.206)	(0.221)	(0.214)	(0.274)
$\ln(Recommended_i)$	0.040**	-0.025	0.107***	0.107***	-0.020	0.208***
	(0.016)	(0.017)	(0.023)	(0.022)	(0.025)	(0.027)
%NSF $Funding_i$	-0.088*	-0.066	-0.104*	-0.011	0.004	-0.006
	(0.050)	(0.061)	(0.062)	(0.084)	(0.084)	(0.097)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field FE	Yes	Yes	Yes			
Field × Year FE	Yes	Yes	Yes			
Dep Var Mean	-0.12	-0.48	0.20	1.64	0.94	2.10
R-Sq	0.29	0.32	0.33	0.05	0.03	0.10
Observations	681	681	681	618	618	618

Notes: The unit of observation is a grant. In Column (1), the outcome is the average SNIP value of papers associated with a grant, logged. In Column (2), the outcome is the minimum SNIP value of papers associated with a grant, logged. In Column (3), the outcome is the maximum SNIP value of papers associated with a grant, logged. In Column (4), the outcome is the average number of citations associated with a grant, logged. In Column (5), the outcome is the number of citations received by the least cited paper associated with a grant, logged. In Column (6), the outcome is the number of citations received by the most cited paper associated with a grant, logged. Citations associated with each paper are normalized to account for differences across fields and publication year. $\%Allocated$ is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. $\ln(Recommended)$ is the log of the number of SUs recommended by the XRAC. $\%NSF\ Funding$ is the share of funding coming from the NSF. Standard errors are clustered at the grant start year × field level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

APPENDIX

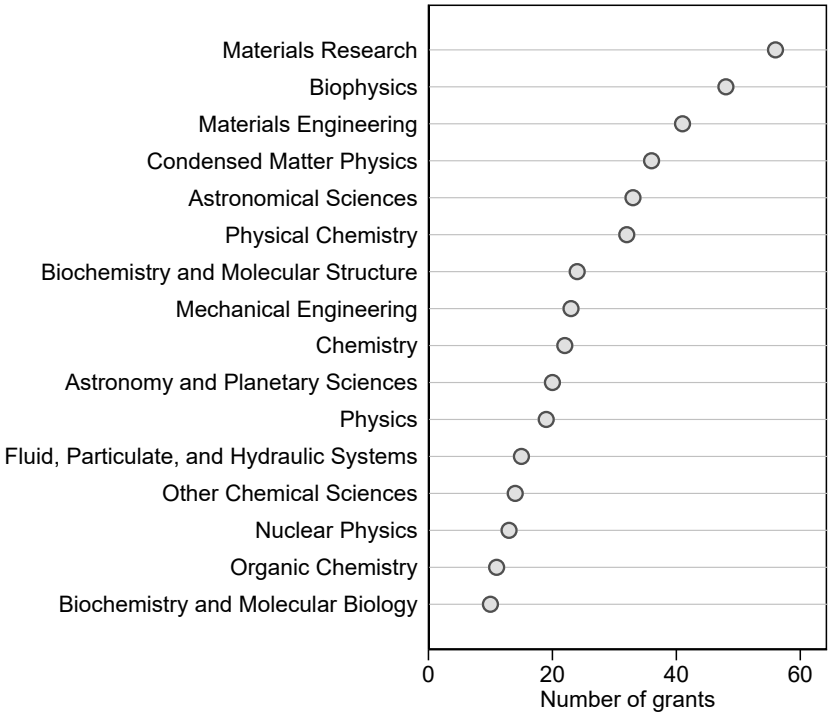
A Figures

Figure A1: XSEDE Resource Allocation Process



Notes: This figure summarizes the resource allocation process of XSEDE.

Figure A2: Most common fields of science (top 15)



Notes: This figure reports the number of grants per field of science for the most common 15 fields of science.

B Tables: Data construction and Empirical strategy

Table A1: Ten Most Used XSEDE Supercomputers

Supercomputer Name	SU (Millions)
TACC Dell/Intel Knights Landing (Stampede2)	27,454
SDSC Dell Cluster (Comet)	10,327
SDSC Expanse CPU	9,306
PSC Bridges-2 (Bridges-2)	8,493
TACC Dell PowerEdge (Stampede)	6,320
PSC Regular Memory (Bridges)	6,079
Purdue Anvil CPU	3,674
NCSA Delta GPU (Delta GPU)	2,340
IU/TACC (Jetstream)	2,264
PSC Bridges-2 GPU (Bridges-2 GPU)	1,965

Notes: This table shows the 10 most used supercomputers over the period 2015-2022. Usage is reported in Service Units (SU) which allows comparison across supercomputers.

Table A2: Sample construction

Sample	Observations
Grants	1,062
Grants with identifiable PIs	950
Grants with a 1:1 correspondence with their PIs	949
Grants with at least one associated paper	719
Grants with a ratio ≤ 1	710

Notes: This table shows the successive steps that lead to the construction of the final sample with 710 observations.

Table A3: Balance test

	DV = %Allocated				
	(1)	(2)	(3)	(4)	(5)
Junior PI	-0.002				-0.011
	(0.017)				(0.016)
R1 Univ		-0.021			-0.009
		(0.020)			(0.020)
Nb papers pre-XSEDE			-0.000		-0.000
			(0.000)		(0.000)
Nb cites pre-XSEDE				-0.000	-0.000
				(0.000)	(0.000)
%NSF $Funding_i$	0.100***	0.096***	0.101***	0.100***	0.100***
	(0.017)	(0.017)	(0.018)	(0.018)	(0.018)
ln($Recommended_i$)	-0.026***	-0.027***	-0.025***	-0.025***	-0.024***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Field FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes	Yes	Yes
Dep Var Mean	0.85	0.85	0.85	0.85	0.85
R-Sq	0.51	0.49	0.52	0.52	0.52
Observations	686	710	675	675	675

Notes: The unit of observation is a researcher. *%Allocated* is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. ln(*Recommended*) is the log of the number of SUs recommended by the XRAC. *Junior* is an indicator equal to 1 the PI has less than 8 years of experience. *R1* is an indicator equal to 1 if the PI is affiliated with an R1 institution. *Nb papers pre-XSEDE* is the number of publications published by the PI before entering XSEDE. *Nb cites pre-XSEDE* is the number of citations received by the PI for publications published before entering XSEDE. All columns control grant start year fixed effects, field fixed effects, and grant start year × field fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Over-requesting test

	DV = $\%Allocated$
	(1)
$\%Over_request$	-0.001
	(0.001)
%NSF $Funding_i$	0.095***
	(0.017)
ln($Recommended_i$)	-0.028***
	(0.005)
Year FE	Yes
Field FE	Yes
Field × Year FE	Yes
Dep Var Mean	0.85
R-Sq	0.49
Observations	710

Notes: The unit of observation is a grant. For each grant, $\%Over_request$ is equal to $(Requested_i - Recommended_i)/Recommended_i$, which we use as a proxy for the share of over-requested resources. $\%NSF\ Funding$ is the share of funding coming from the NSF. ln(*Recommended*) is the log of the number of SUs recommended by the XRAC. We include grant start year fixed effects, grant field fixed effects and grant start year × field fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: Matched vs Unmatched Grants

	DV = 1{In-sample}
	(1)
$\%Allocated_i$	-0.002
	(0.001)
%NSF $Funding_i$	0.077**
	(0.034)
$\ln(Recommended_i)$	0.022*
	(0.013)
Field FE	Yes
Year FE	Yes
Field × Year FE	Yes
Dep Var Mean	0.67
R-Sq	0.38
Observations	1062

Notes: The unit of observation is a grant. The dependent variable is an indicator equal to 1 if the grant is part of our sample of 710 grants and 0 otherwise. $\%Allocated$ is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. $\ln(Recommended)$ is the log of the number of SUs recommended by the XRAC. $\%NSF\ Funding$ is the share of funding coming from the NSF. We control for grant start year fixed effects, grant field fixed effects grant start year × field fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

C Tables: Robustness

Table A6: Relabeling Test

	ln(non-XSEDE papers)					Num. non-XSEDE Papers
	(1)	(2)	(3)	(4)	(5)	(6)
$\%Allocated_i$	-0.643*	-0.248	-0.258	0.082	-0.018	-0.472
	(0.346)	(0.345)	(0.354)	(0.402)	(0.514)	(0.535)
$\ln(Recommended_i)$		0.256***	0.255***	0.246***	0.277***	0.223***
		(0.033)	(0.033)	(0.043)	(0.052)	(0.060)
%NSF $Funding_i$			0.016	-0.085	-0.225	-0.027
			(0.111)	(0.119)	(0.140)	(0.182)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field FE				Yes	Yes	Yes
Field × Year FE					Yes	Yes
Specification	OLS	OLS	OLS	OLS	OLS	Poisson
Dep Var Mean	2.79	2.79	2.79	2.82	2.89	36.75
R-Sq	0.03	0.11	0.11	0.23	0.36	
Observations	654	654	654	654	654	654

Notes: For each grant, we calculate the number of non-XSEDE papers produced while receiving XSEDE resources. In Columns (1)-(5), the outcome is the log number of non-XSEDE papers associated with a grant. Column (6) shows a Poisson specification where the outcome is the number of non-XSEDE papers associated with a grant. $\%Allocated$ is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. $\ln(Recommended)$ is the log of the number of SUs recommended by the XRAC. $\%NSF\ Funding$ is the share of funding coming from the NSF. All columns include grant start year fixed effects. Columns (4) to (6) add grant field fixed effects. Columns (5) and (6) add grant start year × grant field fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: Less traveled research avenues - Robustness

	DV = 1{Use ≥ Median}		
	(1)	(2)	(3)
Paper w/i grant:	Avg. Paper	Least Conv.	Most Conv.
$\%Allocated_i$	-0.282**	-0.578***	-0.134
	(0.110)	(0.149)	(0.120)
$\ln(Recommended_i)$	0.031*	-0.059***	0.057***
	(0.017)	(0.019)	(0.017)
%NSF $Funding_i$	-0.011	0.024	0.038
	(0.044)	(0.052)	(0.048)
Year FE	Yes	Yes	Yes
Field FE	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes
Dep Var Mean	0.62	0.58	0.62
R-Sq	0.51	0.42	0.49
Observations	653	653	653

Notes: The unit of observation is a grant. We first calculate for each paper the average number of times the concepts it entails have been used by other PIs in the same field. We then calculate the mean of this measure at the grant level (mean use of papers within a grant), the minimum of this measure at the grant level (use of concepts by the least conventional paper), and the maximum of this measure at the grant level (use of concepts by the most conventional paper). For each of these grant-level variables, we calculate the median across all grants and create an indicator equal to 1 if the value of the variable for the focal grant is above this median. $\%Allocated$ is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. $\ln(Recommended)$ is the log of the number of SUs recommended by the XRAC. $\%NSF\ Funding$ is the share of funding coming from the NSF. All columns include grant start year fixed effects, grant field fixed effects and grant start year × grant field fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: Less traveled research avenues - Trend over time

	DV = log(Use Prior)			DV = log(Use After)		
	(1)	(2)	(3)	(4)	(5)	(6)
Paper w/i grant:	Avg. Paper	Least Conv.	Most Conv.	Avg. Paper	Least Conv.	Most Conv.
$\%Allocated_i$	-0.616**	-1.640***	-0.217	-0.759***	-2.281***	-0.285
	(0.295)	(0.464)	(0.317)	(0.283)	(0.474)	(0.293)
$\ln(Recommended_i)$	0.034	-0.120**	0.132***	0.053	-0.034	0.170***
	(0.044)	(0.054)	(0.047)	(0.051)	(0.072)	(0.051)
%NSF $Funding_i$	0.050	-0.018	0.024	-0.055	0.026	-0.099
	(0.101)	(0.145)	(0.121)	(0.132)	(0.146)	(0.146)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Dep Var Mean	-0.38	-1.20	0.10	-0.50	-0.99	0.04
R-Sq	0.68	0.57	0.64	0.69	0.63	0.65
Observations	609	512	609	585	416	585

Notes: The unit of observation is a grant. In Columns (1) to (3), we first calculate for each paper the average number of times the concepts it entails have been used by other PIs in the same field *prior* to their entry into XSEDE. Column (1) takes the mean of this value at the grant level, logged. Column (2) takes the min of this value at the grant level, logged. Column (3) takes the max of this value at the grant level, logged. Columns (4) to (6) repeat this exercise using the average number of times concepts used in a paper have been used by other PIs in the same field *after* their entry into XSEDE. $\%Allocated$ is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. $\ln(Recommended)$ is the log of the number of SUs recommended by the XRAC. $\%NSF\,Funding$ is the share of funding coming from the NSF. All columns include grant start year fixed effects, grant field fixed effects and grant start year × grant field fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: Topics beyond expertise - Robustness

	DV = Share of new concepts (vs prev. 3y)			DV = Share of new concepts (vs prev. 5y)		
	(1)	(2)	(3)	(4)	(5)	(6)
Paper w/i grant:	Avg. Paper	Lowest Share New	Highest Share New	Avg. Paper	Lowest Share New	Highest Share New
$\%Allocated_i$	0.101*	-0.010	0.147***	0.091	-0.028	0.145***
	(0.060)	(0.108)	(0.049)	(0.063)	(0.107)	(0.052)
$\ln(Recommended_i)$	-0.001	-0.042***	0.016**	-0.004	-0.045***	0.016**
	(0.008)	(0.014)	(0.007)	(0.008)	(0.014)	(0.007)
%NSF $Funding_i$	-0.000	0.011	-0.016	-0.001	0.013	-0.017
	(0.025)	(0.041)	(0.023)	(0.024)	(0.040)	(0.022)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Dep Var Mean	0.76	0.61	0.87	0.74	0.59	0.86
R-Sq	0.45	0.37	0.43	0.45	0.37	0.44
Observations	608	608	608	617	617	617

Notes: The unit of observation is a grant. For each paper, we calculate the share of 'new-to-the-researcher' concepts, i.e., concepts that have not been previously used by the PI. Columns (1) to (3) compare XSEDE papers with papers written in the 3 years that precede XSEDE entry. Columns (4) to (6) compare XSEDE papers with papers written in the 5 years that precede XSEDE entry. Columns (1) and (4) take the average of this measure at the grant level. Columns (2) and (5) take the min of this measure at the grant level. Columns (3) and (6) take the max of this measure at the grant level. $\%Allocated$ is the share of resources allocated to a grant relative to the amount of resources recommended by the XRAC. $\%NSF\,Funding$ is the share of funding coming from the NSF. All models include grant start year fixed effects, field of science fixed effects and grant start year fixed effects × field of science fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

D Tables: Citations and Research Direction

Table A10: Normalized citations and Type of science

	log(Norm. Cites)			
	(1)	(2)	(3)	(4)
log(Use)	0.080**			
	(0.033)			
Concept Year		-0.020**		
		(0.008)		
Share New (2y)			-0.663***	
			(0.171)	
log(Unique Concepts)				0.234***
				(0.074)
ln(*Recommended*$_i$)	0.078***	0.095***	0.094***	0.100***
	(0.027)	(0.024)	(0.025)	(0.024)
%NSF *Funding*$_i$	-0.021	-0.022	0.003	-0.061
	(0.086)	(0.080)	(0.080)	(0.082)
Year FE	Yes	Yes	Yes	Yes
Dep Var Mean	1.67	1.65	1.66	1.65
R-Sq	0.06	0.06	0.12	0.06
Observations	542	585	524	575

Notes: The unit of observation is a grant. The outcome is the average number of citations associated with a grant, logged. Citations associated with each paper are normalized to account for differences across fields and publication year. To calculate *Log(Use)*, we measure the average number of times concepts used in a paper are used by other PIs and take the mean of this value at the grant level, logged. To calculate *Concept Year*, we calculate for each paper associated with a grant the average year we observe the concepts it entails for the first time and average this measure at the grant level. To calculate *Share New (2y)*, we calculate for each paper the share of concepts that have not been used by the PI in papers written in the 2 years that precede XSEDE entry and average this measure at the grant level. To calculate *Log(Unique Concepts)*, we calculate for each paper the number of unique concepts it entails and average this measure at the grant level, logged. ln(*Recommended*) is the log of the number of SUs recommended by the XRAC. *%NSF Funding* is the share of funding coming from the NSF. All columns control for grant start year fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

E Tables: Main analyses using allocation levels

Table A11: Quantity of papers

	DV = ln(papers)					DV = Num. Papers
	(1)	(2)	(3)	(4)	(5)	(6)
ln(*Allocated*$_i$)	0.275***	0.683***	0.669***	0.537***	0.730***	0.583**
	(0.022)	(0.136)	(0.141)	(0.150)	(0.163)	(0.240)
ln(*Recommended*$_i$)		-0.410***	-0.398***	-0.258*	-0.451***	-0.271
		(0.135)	(0.139)	(0.149)	(0.162)	(0.225)
%NSF *Funding*$_i$			0.034	0.026	-0.076	-0.048
			(0.074)	(0.084)	(0.103)	(0.151)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field FE				Yes	Yes	Yes
Field × Year FE					Yes	Yes
Specification	OLS	OLS	OLS	OLS	OLS	Poisson
Dep Var Mean	1.29	1.29	1.29	1.31	1.35	6.58
R-Sq	0.23	0.24	0.24	0.36	0.46	
Observations	710	710	710	710	710	710

Notes: The unit of observation is a grant. The outcome is the log number of papers associated with a grant. ln(*Allocated*) is the log of the number of SUs allocated. ln(*Recommended*) is the log of the number of SUs recommended by the XRAC. *%NSF Funding* is the share of funding coming from the NSF. We include grant start year fixed effects, grant field fixed effects and grant start year × grant field fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A12: Less traveled research avenues

	DV = log(Use)			DV = Concept Vintage		
	(1)	(2)	(3)	(4)	(5)	(6)
Paper w/i grant:	Avg. Paper	Least Conv.	Most Conv.	Avg. Paper	Oldest	Newest
ln($Allocated_i$)	-0.413**	-1.226***	-0.167	0.967	-1.001	3.285**
	(0.192)	(0.311)	(0.204)	(0.998)	(1.342)	(1.296)
ln($Recommended_i$)	0.472**	1.102***	0.324	-1.332	-0.266	-2.869**
	(0.201)	(0.315)	(0.215)	(1.014)	(1.353)	(1.291)
%NSF $Funding_i$	-0.026	-0.083	-0.053	-0.632	-0.632	-1.027
	(0.120)	(0.145)	(0.137)	(0.473)	(0.685)	(0.673)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Dep Var Mean	-0.41	-1.27	0.07	2003.85	2000.22	2007.54
R-Sq	0.67	0.55	0.63	0.44	0.44	0.39
Observations	613	529	613	665	665	665

Notes: The unit of observation is a grant. In Columns (1) to (3), we first calculate for each paper the average number of times the concepts it entails have been used by other PIs in the same field. Column (1) takes the mean of this value at the grant level, logged. Column (2) takes the min of this value at the grant level, logged. Column (3) takes the max of this value at the grant level, logged. In Columns (4) to (6), we calculate for each concept associated with a paper the first year we see it appearing and take the average at the paper level. Column (4) averages this measure at the grant level. Column (5) takes the min of this measure at the grant level. Column (6) takes the max of this measure at the grant level. ln(*Recommended*) is the log of the number of SUs recommended by the XRAC. $\%NSF\,Funding$ is the share of funding coming from the NSF. All columns include grant start year fixed effects, grant field fixed effects and grant start year × grant field fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A13: Topics beyond expertise

	DV = Share of new concepts (vs prev. 2y)		
	(1)	(2)	(3)
Paper w/i grant:	Avg. Paper	Lowest Share New	Highest Share New
$\ln(Allocated_i)$	0.048	-0.040	0.083**
	(0.038)	(0.063)	(0.034)
$\ln(Recommended_i)$	-0.050	-0.003	-0.069**
	(0.039)	(0.064)	(0.035)
%NSF $Funding_i$	0.007	0.022	-0.006
	(0.026)	(0.044)	(0.023)
Year FE	Yes	Yes	Yes
Field FE	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes
Dep Var Mean	0.78	0.63	0.88
R-Sq	0.45	0.37	0.42
Observations	600	600	600

Notes: The unit of observation is a grant. For each paper, we calculate the share of 'new-to-the-researcher' concepts, i.e., concepts that have not been previously used by the PI in the 2 years that precede XSEDE entry. Column (1) takes the average of this measure at the grant level. Column (2) takes the min of this measure at the grant level. Column (3) takes the max of this measure at the grant level. ln(*Allocated*) is the log of the number of SUs allocated. ln(*Recommended*) is the log of the number of SUs recommended by the XRAC. $\%NSF\ Funding$ is the share of funding coming from the NSF. All models include grant start year fixed effects, field of science fixed effects and grant start year fixed effects × field of science fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A14: Topics beyond expertise - alternative time windows

	DV = Share of new concepts (vs prev. 3y)			DV = Share of new concepts (vs prev. 5y)		
	(1)	(2)	(3)	(4)	(5)	(6)
Paper w/i grant:	Avg. Paper	Lowest Share New	Highest Share New	Avg. Paper	Lowest Share New	Highest Share New
ln(*Allocated*$_i$)	0.064	-0.010	0.096***	0.059	-0.020	0.097***
	(0.040)	(0.072)	(0.033)	(0.042)	(0.072)	(0.034)
ln(*Recommended*$_i$)	-0.066	-0.032	-0.081**	-0.064	-0.025	-0.081**
	(0.040)	(0.072)	(0.033)	(0.042)	(0.071)	(0.035)
%NSF *Funding*$_i$	0.000	0.011	-0.016	-0.001	0.014	-0.016
	(0.025)	(0.041)	(0.023)	(0.024)	(0.040)	(0.022)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Dep Var Mean	0.76	0.61	0.87	0.74	0.59	0.86
R-Sq	0.45	0.37	0.43	0.45	0.37	0.44
Observations	608	608	608	617	617	617

Notes: The unit of observation is a grant. For each paper, we calculate the share of 'new-to-the-researcher' concepts, i.e., concepts that have not been previously used by the PI in the years that precede XSEDE entry. Columns (1) to (3) compare XSEDE papers with papers written in the 3 years that precede XSEDE entry. Columns (4) to (6) compare XSEDE papers with papers written in the 5 years that precede XSEDE entry. Columns (1) and (4) take the average of this measure at the grant level. Columns (2) and (5) take the min of this measure at the grant level. Columns (3) and (6) take the max of this measure at the grant level. ln(*Allocated*) is the log of the number of SUs allocated. ln(*Recommended*) is the log of the number of SUs recommended by the XRAC. $\%NSF\ Funding$ is the share of funding coming from the NSF. All models include grant start year fixed effects, field of science fixed effects and grant start year fixed effects × field of science fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A15: Research Breadth

	DV = log(Unique Concepts)		
	(1)	(2)	(3)
Paper w/i grant:	Avg. Paper	Narrowest Paper	Broadest Paper
ln(*Allocated*$_i$)	0.267**	-0.190	0.435***
	(0.114)	(0.199)	(0.125)
ln(*Recommended*$_i$)	-0.279**	-0.068	-0.394***
	(0.113)	(0.196)	(0.125)
%NSF *Funding*$_i$	0.042	0.157	0.020
	(0.046)	(0.124)	(0.054)
Year FE	Yes	Yes	Yes
Field FE	Yes	Yes	Yes
Field × Year FE	Yes	Yes	Yes
Dep Var Mean	3.09	2.32	3.37
R-Sq	0.35	0.41	0.36
Observations	653	653	653

Notes: The unit of observation is a grant. We calculate for each paper the number of unique concepts it entails. Column (1) averages this measure at the grant level, logged. Column (2) takes the min of this measure at the grant level, logged. Column (3) takes the max of this measure at the grant level, logged. ln(*Allocated*) is the log of the number of SUs allocated. ln(*Recommended*) is the log of the number of SUs recommended by the XRAC. *%NSF Funding* is the share of funding coming from the NSF. All columns include grant start year fixed effects, field of science fixed effects and grant start year fixed effects × field of science fixed effects. Standard errors are clustered at the grant start year × field level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

F Figures: Non-parametric specification

To construct the following figures, we bin the *% Allocated* variable into 5 quantiles: below the 10^{th} percentile, between the 10^{th} percentile and the 25^{th} percentile, between the 25^{th} percentile and the 50^{th} percentile, between the 50^{th} percentile and the 75^{th} percentile and above the 75^{th} percentile. We run similar regressions as Equation 1 using these indicators as main variables of interest and plot the coefficient estimates we obtain. The dependent variables are the same as described in the text.

Figure A3: Papers

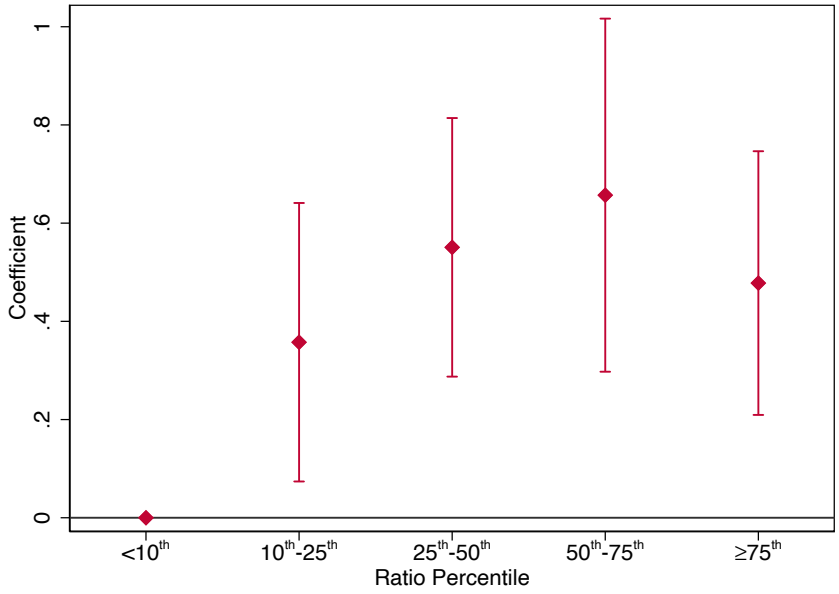


Figure A4: Use of concepts by others

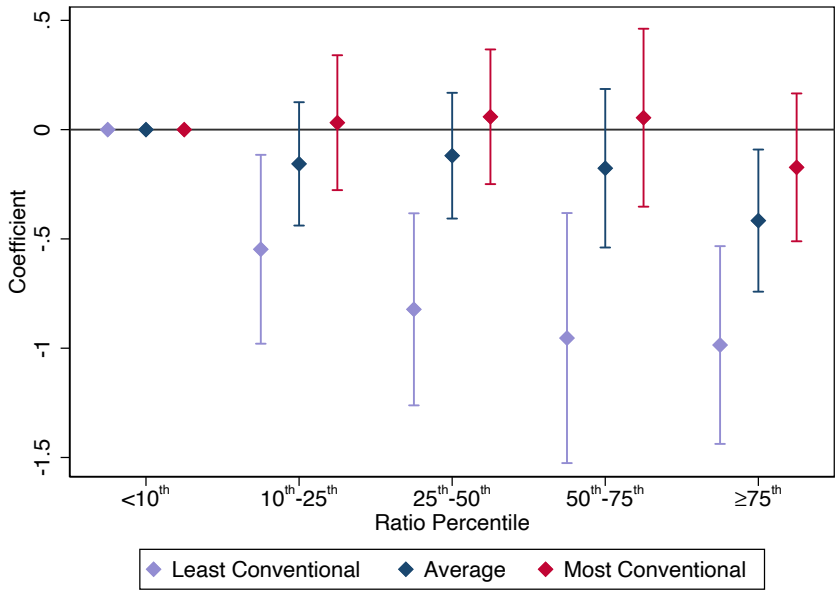


Figure A5: Concept Vintage

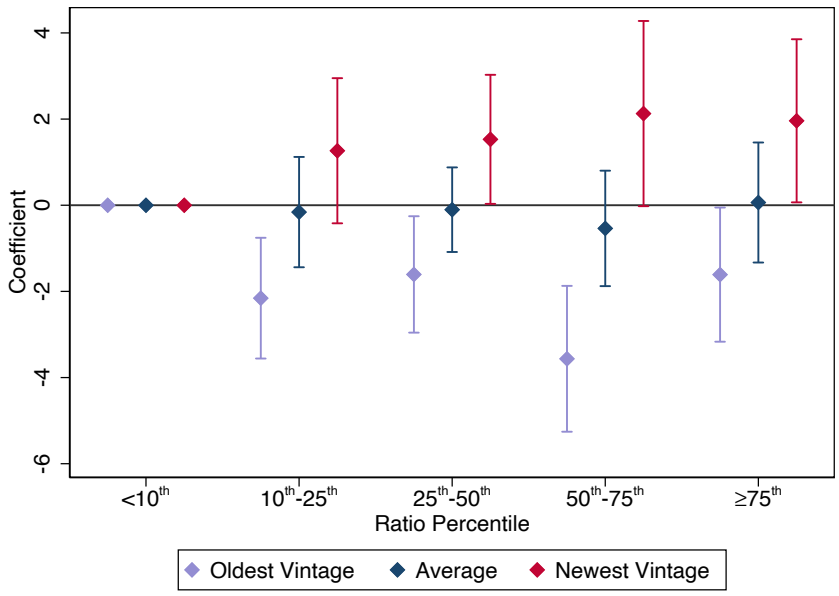


Figure A6: Share of new concepts

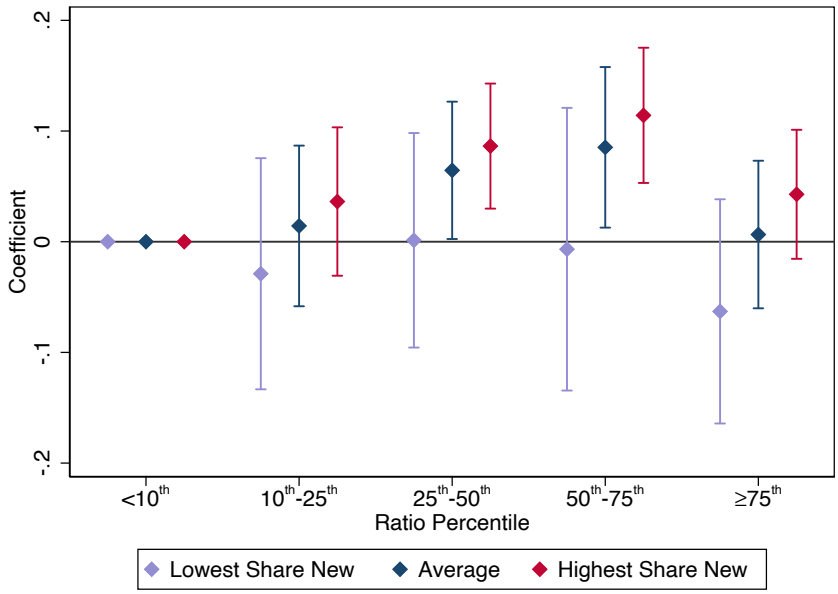


Figure A7: Research breadth

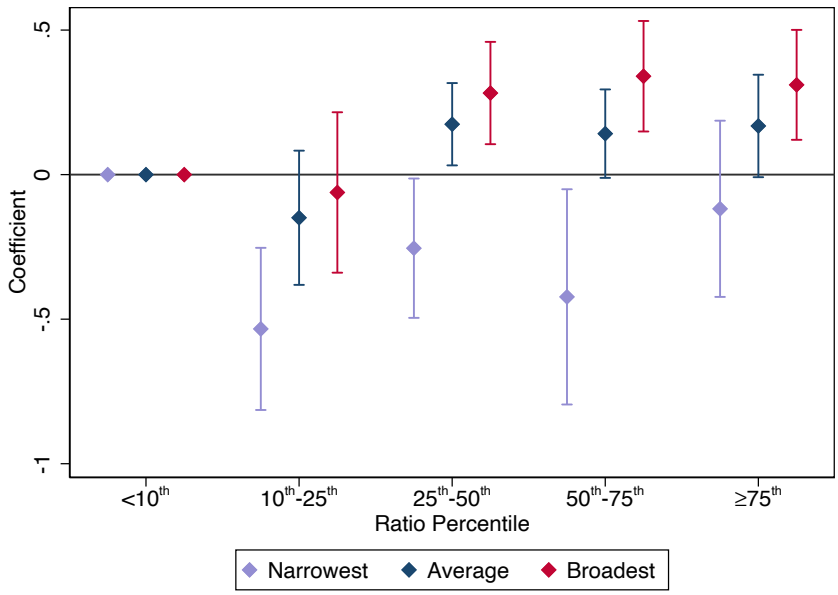


Figure A8: SNIP

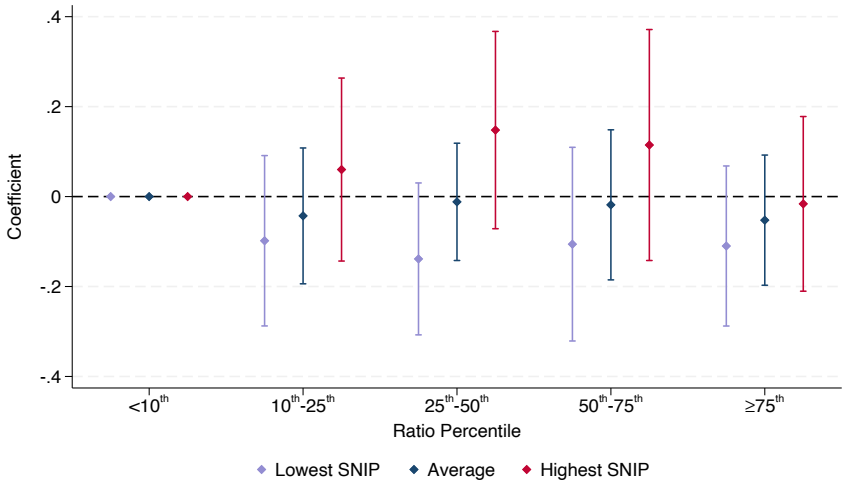
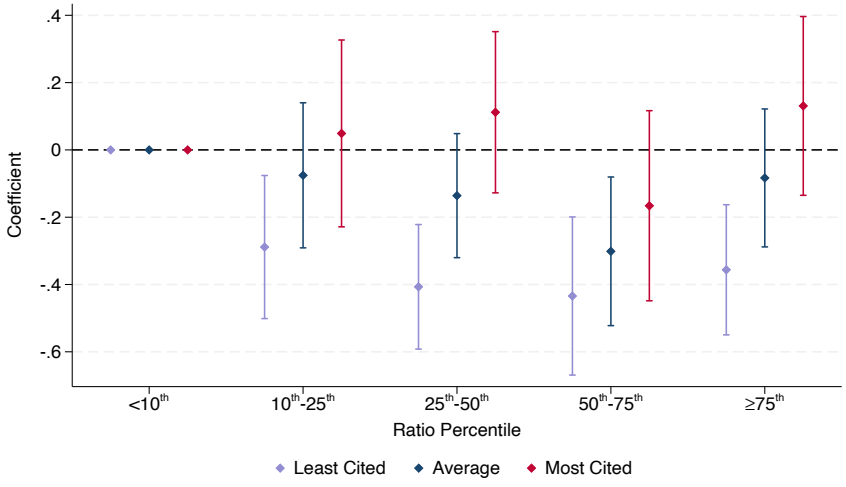


Figure A9: Citations



G Proposal Example

This is an example of a proposal in the field of Biophysics.[30]

G.1 Summary of Proposed Research

We are requesting continued use of XSEDE resources to support four independent projects that each utilize conventional and enhanced sampling molecular dynamics (MD) simulations. Specifically, these projects are:

1. The nucleosome core particle (NCP) consists of ∼147 base pairs of DNA wrapped around a histone octamer and is the fundamental unit of DNA packaging in eukaryotes. The substitution of variant proteins to the octamer has been linked to epigentic regulation in vivo, such as the increased presence of H2A.B in transcriptionally active regions of DNA. While several studies have highlighted potential mechanisms, a fully atomistic understanding of H2A.B-regulated NCP destabilization has not been described. Here, we propose a combination of conventional and free energy MD simulations to study the mechanisms of epigenetic activation by H2A.B substitution.

2. The chromatosome is an extension of the NCP which binds an additional ∼20 base pairs of DNA due to the presence of a linker histone. Experiments have shown that linker histones can bind to DNA in either an on- or off-dyad locations, the selection of which is dictated by modifications to the linker histone sequence. However, the molecular mechanisms that drive the binding mode, as well as how they are influenced by specific linker histone residues, remains poorly understood. Here, we will perform a series of conventional MD simulations of three different linker histones bound in on- and off-dyad locations to determine how specific residues influence the structure and energetics of the chromatosome.

3. Chromatin fibers consist of repeating NCP units and form the building blocks of chromosomes within the nucleus of eukaryotic cells. Therefore, the structure and dynamics of poly-NCP arrays directly influences several processes, including gene expression, DNA replication, and DNA repair. In this project, we aim to study the relationship between structure, function, and dynamics in poly-NCP

[30] See `https://web.archive.org/web/20210414012547/https://portal.xsede.org/allocations/research#examples` for more details, in-text citations from this proposal, and additional examples.

arrays. More concretely, we will perform fully atomistic MD simulations of an octa-NCP construct in a compact state. Analysis will focus on the stability and conformational dynamics of the chromatin fiber, along with the importance of stacking interactions between adjacent nucleosomes and inter-NCP contacts formed by the histone tails. Furthermore, we will focus on identifying allosteric networks within poly-NCP arrays.

4. Pili are protein structures that protrude from bacterial cell walls and function to enhance their virulence. They are bio-synthesized by the C class sortase proteins, which was the subject of our previous work. Structurally, pili are polymers built by covalently linking various pilin subunits. However, there is little detail in the literature about the molecular structure and dynamics of these building blocks or how their properties are potentially modified due to polymerization. Here, we propose to implement free energy MD simulations on the dimeric form of Streptococcus pneumoniae and Corynebacterium diphtheriae pilins. Simulations will be carried out both in the presence and absence of the experimentally determined intradomain isopeptide bonds that have been proposed to impart mechanical strength to pili.

G.2 Background and Specific Goals

G.2.1 Simulations of the H2A.B Variant Nucleosome (NIH 1R15GM114758)

Eukaryotic cells store their genetic material into chromosomes through a hierarchical packaging scheme whose fundamental unit is the nucleosome core particle (NCP) [1]. The NCP wraps ∼147 base pairs of DNA around a symmetric histone octamer formed by two (H2A-H2B) dimers anking an (H3-H4)2 tetramer [2]. In vivo, NCPs act as epigenetic governors of transcription and alter gene expression through a variety of chemical changes [3, 4]. For example, substitution of canonical histones with a structurally similar but sequentially deviant protein, a histone "variant", can comprehensively modulate gene expression. Our previous work was focused on a transcription repressor, the macroH2A variant, in both homogeneous (both dimers replaced) [5] and heterogeneous (one canonical dimer and one variant dimer) nucleosomes. In contrast, Barr body deficient H2A (H2A.B) is associated with transcriptionally active regions of DNA[6], but little is known about how the structure and dynamics of H2A.B containing NCPs affect their function.

Like other H2A variants, the histone fold of H2A.B is believed to follow the structural motif of three alpha helices connected by short loops and anked by exible N- and C-terminal tails. However, the sequence of this fold is only ∼50% identical to the canonical protein, making H2A.B the most sequentially divergent of the major variants [3]. Indeed, these sequence variations alter important chemical properties of the NCP, such as reducing the total number of lysine residues and eliminating the acidic patch believed to be crucial in internucleosomal interactions. Strikingly, the C-terminal sequence of H2A.B is ∼20 amino acids shorter than the canonical protein, reducing the sequence responsible for forming dimer-tetramer interactions by ∼25%. Furthermore, these C-terminal residues are located near the DNA entry/exit site in canonical H2A. Thus, the lack of C-terminal sequence may destabilize the NCP by degrading both protein-protein and protein-DNA interactions and result in an open, more accessible structure [79].

In the simulations proposed here, we will comprehensively study the eects of H2A.B substitution on NCP dynamics and stability through a combination of conventional and free energy MD simulations. Results from conventional simulations will be compared against our previous work to determine the differences in dynamics between H2A.B and canonical nucleosomes in the hundreds of nanoseconds timescale, and solution ensembles will be compared against experimental scattering profiles [7, 8]. Furthermore, free energy calculations will be conducted for both the canonical and H2A.B NCPs to quantify the disparity in DNA binding capabilities. The success of these free energy calculations will serve to guide unwinding simulations on future NCP systems, such as other variant NCPs, the chromatosome, and post-translationally modied NCPs. These simulations, along with our previous work, will provide further understanding of the interaction between structure, dynamics, and function in variant-containing nucleosomes.

G.2.2 Simulations of the Chromatosome (NSF 1552743)

The chromatosome is an extension of the NCP that contains the same structural foundations with an additional ∼20 base pairs of DNA (∼167 total) accompanied by a highly-conserved linker histone. Colloquially known as histone H1, this nuclear protein plays a crucial role in the condensation of nucleosome chains into higher order structures [10], as well as other cellular functions [11] such as gene expression [12, 13], heterochromatin genetic activity [14], and cell differentiation [15,16], among many others [1719]. Linker histones have also been found to rapidly exchange locations throughout the chromosome [20, 21]. In avian species,

H1 is replaced by the H5 isoform.

Although H1 and H5 have highly similar primary and tertiary structures, recent experiments have shown that they exhibit different binding modes on the nucleosome. Histone H5 binds centered on the dyad axis in the "on-dyad" location, whereas histone H1 binds in a skewed mode relative to the dyad axis, or off-dyad. In a recent study, Zhou et al examined the binding modes of wild type globular histone H1 (GH1), H5 (GH5), and a GH5 penta-mutant that had five residues switched from their GH5 to GH1 sequences (GH5 pMu) [22]. Results showed that these small number of mutations were able to shift the binding state from the H5 on-dyad location to the H1 off-dyad location. This suggests that these five residues may be the key determinants for linker histone binding modes on chromatosomes. However, the molecular mechanisms that orchestrate this conformational shift, the role of specific linker histone residues, and how residue modifications affects this process is poorly understood. Moreover, the extent to which these phenomenon alter greater chromatin dynamics is also unclear.

Here, we will study the atomic scale mechanisms that dictate linker histone binding modes through a series of conventional MD simulations of different chromatosomes containing GH1, GH5, and GH5 pMu. We hypothesize that each linker histone is capable of binding in both the on- and off-dyad locations, and that differences in the histones shift the free energies of the states. Therefore, we will initiate simulations of each system from both the on- and off-dyad locations. Analysis will focus on the structure and energetics of specific contacts that stabilize these two binding modes, with an emphasis on how the mutated residues in GH5 pMu create the experimentally observed shift in binding.

G.2.3 Simulations of the Canonical Octanucleosome (NSF 1552743)

In vitro, nucleosomes form long "beads on a string" like structures in which many nucleosomes are periodically positioned on a long DNA molecule. In its most compact form, poly-NCP arrays condense into structures that are composed of tetra-NCP repeats with DNA that zigzags between the individual units [23]. These tetra-NCP structures likely form the basis of chromatin fibers and thus represent a second level of gene compaction [1, 2426]. What the relationship is between the structure, dynamics, and function of these constructs, and how these properties are modulated by chromatin remodeling factors, are questions of deep scientific interest that have yet to be fully understood.

In previous work on mono-NCPs, we showed that chromatin remodeling factors influence not only the NCP structure, but also the dynamic networks that extend throughout the complex [5]. We hypothesize that the regulation of allosteric networks is a general feature of chromatin remodeling factors, and that histone variants and PTMs modulate both intra- and inter-NCP dynamics as a method of influencing chromatin stability far from their physical location. If true, these networks would likely propagate through the inter-nucleosomal DNA and histone tails, and they would link locations that affect nucleosome stability, such as the DNA entry and exit points, as well as the locations most heavily influenced by remodeling factors, such as the L1 loops and the acidic patch.

In this project, we will begin to test the hypothesis that allosteric networks link regions between nucleosomes that are known to be important for chromatin fiber regulation. To do this, we will perform a set of conventional MD simulations of an octa-NCP system. Analysis will focus on determining the strength and linchpins of inter-NCP and inter-tetraNCP dynamic networks with a combination of methods that we previously applied to mono-NCP systems [5]. In future years, we will build on these results by simulating poly-NCP constructs that include histone variants and PTMs, as well as systems in other, non-compact, conformations. We also aim to use our full atomistic simulations to validate existing Coarse Grained models that have been used to simulate mono-NCPs [27] and their reliability to be used in simulating larger poly-NCP to model realistic chromatin fibers.

G.2.4 Simulations of Pilin Subunits (NIH 1R35GM119647)

Pili, which are elongated protein structures that reside on the cell surface of bacteria, have been shown to increase bacterial virulence. They are formed by sortase C (SrtC) proteins through the polymerizing of pilin sub-units [28]. The main pilin monomer of Streptococcus pneumoniae (RrgB) is made up of four domains as shown in Figure 4. The second and third domains are arranged side to side, while the first and last units lie on top and below these middle units. In the case of Corynebacterium diphtheriae, its main building block (SpaA) is a smaller three-domain protein with domains that are placed on top of each other. Interestingly, each of the domains in both proteins contain isopeptide bonds that have been proposed to provide additional mechanical stability.

Previously, we examined the lid opening process of SrtC, and our results predicted that it likely requires

the action of an external factor to function due to the high calculated free energies required for the lid opening process; a result that was later confirmed by NMR dynamics experiments carried out by our collaborators [29]. We further went on to examine the energetics of pilin sorting signals binding to SrtC (Progress Report, section 1.2). For the present project, we aim to understand the structural dynamics of the Streptococcus pneumoniae and Corynebacterium diphtheriae pilin monomers, explain how dimerization affects overall dynamics, and elucidate the molecular reasons for the intra-domain isopeptide bonds observed in these pilin molecules.

G.3 Proposed Simulations and System Details

Here, we describe the specific details of the simulations we will perform in each project. These simulations will use either NAMD, GROMACS, or AMBER [3032], and GPU acceleration will be used whenever feasible [33]. All simulations with explicit solvent environments will be performed with periodic boundary conditions, with non-bonded interactions truncated at 10 A, and long-range electrostatics treated with the particle mesh Ewald (PME) method [34]. Hydrogen containing bonds will be constrained using SHAKE [35] in AMBER and NAMD, and P-LINCS will restrain similar bonds in GROMACS simulations [36]. All simulations are to be conducted with a 2 fs timestep unless otherwise specified.

G.3.1 Simulations of the H2A.B Nucleosome

To better understand the structure-dynamics-function relationship in H2A.B containing NCPs, we are requesting XSEDE resources for conducting two complementary arrays of simulations. In the first, we will conduct conventional MD (cMD) simulations on the H2A.B NCP. In the second, we will calculate the free energy of unwinding of canonical and H2A.B NCPs via umbrella sampling (US) calculations. In both sets of simulations, the protein and nucleic acid forcefield parameters will come from FF14SB, the most recent AMBER xed point-charge forcefield [37, 38], and ions will be paramaterized according to Joung and Cheatham [39].

G.3.1.1 Conventional Molecular Dynamics

Since no crystal structure of the H2A.B NCP exists, an initial conformation was constructed from the 1.9 A resolution X-ray structure of the canonical NCP (PDB: 1KX5) [40]. The canonical H2A sequence was truncated and mutated to the sequence of H2A.B in VMD and then neutralized by the addition of Cl- ions. The neutralized H2A.B NCP was then solvated in a box of TIP3P water[41] at a salt concentration of 150

Table 1: A summary of the total requested resources for simulations relating to the H2A.B NCP. Conventional simulations have a benchmark speed of 300 SUs/ns on Comet with NAMD, whereas umbrella sampling calculations were benchmarked at 5 SUs/ns on XStream with AMBER.

NCP Type	Simulation Method	System Size (atoms)	Simulation Length	Number of Simulations	Comet SUs	XStream SUs
H2A.B	cMD	241,730	250 ns	3	225,000	———
H2A.B	US	24,509	250 ns	36	———	45,000
Canonical	US	25,172	250 ns	36	———	45,000
				Total:	**225,000**	**90,000**

mM NaCl, and the box was extended 10 A from the solute, yielding a system size of 241,730 atoms.

Simulations will be performed in the NPT ensemble at 300 K and 1 atm using the NAMD engine and a 2 fs timestep. Long range correlations, DNA-protein interactions and binding energies, and dimer-tetramer interactions will be analyzed from the resulting trajectories and will be compared against our previous results for the canonical NCP[5]. Furthermore, the H2A.B solution of states will be compared to experimental scattering data [7, 8] using our recently developed in-house protocol (see Progress Report, Section 1.4).

To determine the amount of necessary simulation time, we direct to our previous work on canonical and macroH2A variant NCPs in which our simulations required 50-60 ns in order for the extended histone tails to equilibrate and contract [5]. Therefore, in order to obtain ∼200 ns of production data, a total of 250 ns per simulation must be conducted. Three separate simulations will be employed to ensure meaningful statistics in observed dynamic behavior. Thus, a total of 750 ns of cMD will be conducted for the H2A.B NCP. To conduct these simulations, we request 225,000 SUs on Comet, in accordance with our benchmarking data. Additionally, we are requesting 0.42 TB of storage space on Data Oasis for this portion of the project.

G.3.1.2 Umbrella Sampling of DNA Unwinding

To determine the simulation protocol for our umbrella sampling calculations, we reference a recent coarse-grained simulation of DNA unwinding from the standard NCP [42]. The reported free energy profile suggests that the unwrapping of DNA occurs in an asymmetric pathway in which only one end of DNA unwinds, in agreement with other studies [43, 44]. To generate their PMF, Zhang et al. coupled harmonic restraints to the radius of gyration (Rg) of all DNA base pairs over 12 windows spanning a range of 45-72.5 A, with centers evenly spaced every 2.5 A. However, sufficient sampling for a similar protocol over such a wide range

of Rg values at the atomistic level is currently unfeasible. Instead, we will utilize the asymmetric unwinding behavior to conduct two separate umbrella sampling calculations per NCP system, one for each DNA end. The reaction coordinate will be defined by the Rg of 73 base pairs extending from the dyad axis in either direction, and harmonic biases will be introduced to 18 windows spread over a range of Rg = 42.5-59.5 A. Note that we are increasing the number of windows relative to the study done by Zhang et al., as we expect that our all-atom model will have smaller fluctuations, and thus require a higher density of windows, then their course grain model. Probability densities from all windows will be converted into a 1-dimensional potential of mean force (PMF) using WHAM [45].

The accuracy of any free energy calculation relies heavily on the convergence of each simulation window. Therefore, we propose that each window will be run for ∼250 ns, in accordance to our cMD protocol. To increase computational performance, we will employ a hydrogen mass repartitioning protocol to increase our timestep two-fold (Dt = 4 fs) [46], and we will treat the solvent environment implicitly, which has been shown to provide good interpretations of nucleosome dynamics [47]. In doing so, we will reduce each system to ∼25,000 atoms to be simulated at 300 K in the NVT ensemble with all pair-wise interactions calculated directly.

We will utilize the effcient GPU performance of the recently developed pmemd.CUDA (v16) for simulating each window. From our benchmarks on XStream, simulating the proposed windows will require 45,000 SUs per NCP type, so we request a total of 90,000 XStream SUs to conduct our umbrella sampling calculations. Furthermore, we request 0.52 TB of storage space on Data Oasis for this portion of the project.

G.3.2 Simulations of the Chromatosome

To determine the structural and energetic factors that dictate the binding states of linker histones, conventional MD simulations of chromatosomes with GH1, GH5, and the GH5 pMu linker histones will be conducted. Each simulation will be initiated in both the on- and off-dyad conformations to determine the how seemingly minor structural differences shift the equilibrium of these states. Simulations of the on-dyad state will be initiated from the 4QLC crystal structure of the H5 chromatosome system, with missing loops and tails modelled using Modeller 9.17 [48]. Currently there is no atomic-resolution structure of a chromatosome containing an off-dyad linker histone. Therefore, these simulations will be based on the cryogenic

electron microscopy (cryo-EM) micrograph of a chromatin fiber which contains histone H1 in the off-dyad binding mode (EMDataBank ID emd 2600.ccp4) [49]. Using this electron density map, we will perform molecular dynamics flexible fitting (MDFF) calculations to fit linker histones to the off-dyad position [50]. These simulations will be performed with local resources.

Simulation parameters will be similar to those for our H2A.B simulations, including the use of an NPT ensemble at 300 K and 1 atm with the NAMD engine and a 2 fs timestep. Protein and nucleic acid parameters will be taken from the AMBER ff14SB force field, and TIP3P waters will be used. As in our previously performed simulations of the canonical and macroH2A-containing NCP, as well as the proposed H2A.B simulations, each simulation will be performed for 250 ns three times. Analysis will focus on the complex stability, as well as the binding energetics and dynamics of the different linker histones.

Using current resources allocated by XSEDE, we have already begun - and expect to finish - three simulations of the chromatosome containing GH5 in the on-dyad binding mode. Therefore, we are requesting resources for a total of 15 separate 250 ns simulations to complete three trajectories of the remaining five systems. In accordance to our benchmarking data, we are requesting 1,312,500 Comet hours (see Table 2). Additionally, we are requesting 2.40 TB of storage to store our trajectory files.

Table 2: A summary of the total requested resources for simulations relating to the chromatosome. A benchmark speed of 350 SUs/ns was measured on Comet.

Linker Histone (Conformation)	Simulation Method	System Size (atoms)	Simulation Length (ns)	Number of Simulations	Comet SUs
GH1 (On-Dyad)	cMD	270,000	250	3	262,500
GH1 (Off-Dyad)	cMD	270,000	250	3	262,500
GH5 (Off-Dyad)	cMD	270,000	250	3	262,500
GH5_pMu (On-Dyad)	cMD	270,000	250	3	262,500
GH5_pMu (Off-Dyad)	cMD	270,000	250	3	262,500
				Total:	**1,312,500**

G.3.3 Simulations of the Canonical Octanucleosome

To begin our study of inter-NCP allosteric networks, we are performing cMD simulations of a compact octanucleosome comprised of canonical histones. The original structure for these simulations is based on the 11 A resolution cryo-EM micrograph of a 30 nm chromatin fiber reconstituted on 12 x 177 base pairs (EMDataBank ID emd 2600.ccp4) [49]. Individual NCPs are based on the 1.9 A crystal structure by Davey

et al. [40] (PDB ID: 1KX5). Each tetra-NCP is based on the crystal structure by Schalch et al.[23] (PDB ID: 1ZBB). We use the AMBER14SB force field to describe the proteins and the DNA. The system is solvated with TIP3P water molecules, and counterions have been added to neutralize the system and simulate a 150 mM NaCl environment [37, 51]. The resulting system has ∼2,800,000 atoms. We use Gromacs 5.0.6 to run these simulations. Pressure is controlled using the Parrinello-Rahman barostat [52, 53] while temperature is controlled using the Nose-Hoover thermostat [54, 55].

To map allosteric effects in the NCP, we will utilize multiple techniques. Residue-residue correlations will be calculated by utilizing mutual information methods [56, 57], while allosteric networks will be mapped with the Weighted Implementation of Sub-optimal Pathways approach [58]. In addition, the importance of individual residues to dynamic networks throughout the structure will be computed with edge-betweenness centrality calculations [59]. Finally, we will determine how packing into chromatin fibers alters the dynamics of a single NCP by computing the Kullback-Leibler divergence of dihedral angle populations between these simulations and our previous monoNCP simulations [60].

To estimate the required simulation times for convergence of these calculations, we refer to our previous NCP study where we performed three independent 250 ns simulations for each construct. Given the system size of the octaNCP, dynamics is expected to be slowed by packing. We estimate that simulations twice as long, or 500 ns in length, will be required to achieve an acceptable level of sampling. We have already performed 250 ns of cMD of a canonical octanucleosome. Preliminary results show that simulation times of ∼150 ns are required to reach equilibrium. In addition, it is important that we perform these simulations multiple times to delineate which dynamical properties are robust. Therefore, we intend to perform three 500 ns simulations. Based on the timing in the benchmark section, we therefore request 3,750,000 Stampede hours (see Table 3). In addition, we are requesting 3.96 TB of storage to save our trajectory files (see scaling document).

G.3.4 Umbrella Sampling Molecular Dynamics Simulations of Pilin Subunits

Using in house resources, we implemented three conventional and three accelerated MD simulations of the monomer and dimer forms of RrgB and SpaA. The results suggest that their structures are more dynamic in the absence of the intra-domain isopeptide bonds. In some of the calculations, a kink appears in the dimer

Table 3: Simulations we are requesting XSEDE resources for in the octanucleosome project. A conversion rate of 2997 SUs/ns was measured on Stampede (see Benchmark data).

System	System Size (Atoms)	Simulation Time	Number of Simulations	Stampede SUs
OctaNCP	2,800,000	500 ns	2	3,000,000
OctaNCP	2,800,000	250 ns	1	750,000
			Total:	**3,750,000**

with the angle between monomers approaching approximately 110°. Such structural changes were not seen in the simulations of dimers with an isopeptide bond, which could imply that isopeptide bonds are vital to pili structural integrity. However, it could also be that the calculations implemented for structures with an isopeptide bond are yet to find this kinked conformation. Clarifying this requires the use of free energy calculations. Therefore, we propose to implement umbrella sampling calculations to further extend the aMD and cMD simulations we have already performed and provide detailed molecular level information about the role of isopeptide bonds in pilin molecules.

The octahedral box containing a solvated RrgB dimer is made up of 987,650 particles. The SpaA system has a slightly smaller total number of particles due to its smaller size. To implement the umbrella sampling calculations, a dihedral collective variable will be used. The vertices of this dihedral will be made from the center of mass of groups of atoms in the first and last sub-units of RrgB and SpaA. These calculations will therefore make use of the non-GPU accelerated version of AMBER, since such groupings are yet to be implemented in the GPU accelerated module. The dihedral angle will be varied from 90° to 180° in 3° increments, leading to thirty windows per system. Each window will be run for 100 ns, which means a total time of 3.1 μs per system, and windows will be combined to create a PMF using WHAM [45]. A 100 ns simulation time was chosen since we observed equilibration times of between 20 to 40 ns in the aforementioned cMD calculations, which will allow 60 to 80 ns of sampling time for collecting statistics. Calculations will be carried out three times for RrgB and SpaA dimers in the presence and absence of isopeptide bonds, yielding a total of 12.0 μs simulation time for this project. Therefore, in accordance with our benchmarking data, we request 744,000 Comet SUs for our free energy calculations of pilin structures. We are also requesting 14.9 TB of storage for these calculations.

Table 4: Simulations we are requesting XSEDE resources for umbrella sampling of pilin proteins. A conversion rate of 60 SUs/ns was used for the simulations (see Benchmark data).

System	Number of Simulations	Length	Comet SUs
RrgB Dimer	100 ns × 31 windows	3.1 µs	186,000
RrgB Dimer with Iso	100 ns × 31 windows	3.1 µs	186,000
SpaA Dimer	100 ns × 31 windows	3.1 µs	186,000
SpaA Dimer with Iso	100 ns × 31 windows	3.1 µs	186,000
		Total:	**744,000**

H XSEDE Review & Reconciliation Process

This section provides additional details about the XSEDE Review and Reconciliation process, through which recommendations are made and then reduced to final allocations. In addition to conversations with individuals at XSEDE, much of the following (including excerpts) comes from the 2021 v1.1 version of *XSEDE Allocations Practices and Procedures* and the 2019 v1.1 version of the *XRAC Reviewer Manual*.[31]

H.1 Review Criteria:

The excerpt below from the *XRAC Reviewer Manual* summarizes the criteria reviewers should consider when making recommendations from proposals:

1. *Appropriateness of Methodology: Does the request describe appropriate tools, methods, and approaches for addressing the research objectives? These methodologies may be community codes or models, data analysis methods, or algorithmic formulations expressed in user-developed scripts or tools.*
2. *Appropriateness of Research Plan: Does the request describe necessary and sufficient computational experiments to answer the research questions posed? In some cases, the research plan may be more reasonably expressed as estimates of resource use, supported by past or early experience. Serious concerns about the research plan will be documented in reviews and may lead to reduced allocation awards.*
3. *Efficient Use of Resources: Has the request identified appropriate resources for undertaking the research plan with the proposed methodology? And will those resources be used as efficiently as is reasonably possible and in accordance with the recommended use guidelines of those resources? Is the proposed work better suited or equally well suited to other resources available for allocation?*

XRAC reviewers may also consider these "additional considerations:"

[31] The reviewer manual can be found at `https://www.ideals.illinois.edu/items/111788/bitstreams/366207/data.pdf`. The practices & procedures document has been removed from the internet but can be shared by the authors upon request.

Intellectual Merit: Reviewers also consider whether the work has any supporting grants for which the science has been merit-reviewed. If so, is the allocation request consistent with the objectives of the supporting grants as described in the submission and is the scale of resource use commensurate with the level of support? If no such grants are identified, the reviewers must assess the intellectual merit of the work. This assessment of the intellectual merit will factor into their overall recommendation.

In our context, the intellectual merit consideration serves to put all requests on a "level playing field." The vast majority of grants are supported by government agencies and do not warrant review of their intellectual merit.

Prior results: Research requests typically require some form of prior work or progress to be considered for an allocation. For New requests, demonstrated benchmarking or test runs are typically required on the requested resource in almost all cases; in some instances (such as a request for a new resource for which Startups are not yet available), demonstrated work on an architecturally similar resource may suffice. For Renewal requests, reviewers consider the provided Progress Report as well as utilization levels of prior allocations

The prior review criteria for renewal requests ensures researchers are using their existing allocations (i.e., are not wasteful). In conversations with XSEDE directors, this criteria serves primarily to check that prior allocations are indeed "doing something." In particular, reviewers may look at if the project has produced *any* publications. If this leads grants with a prior publication to receive a higher recommendation than their *true need*, one might wonder whether this may introduce bias to our quantity (number of papers) analysis. This is not a concern because larger recommendations are reduced more (see Appendix H.2) so this feature of the renewal process would in fact bias our estimates in Table 2 downward, leading us to underestimate the true effect. This can be seen in a simple numerical example. Suppose we observe a project that had already published a paper using a previous XSEDE allocation and the PI requests a renewal. This project is recommended 100 SU and is awarded 50 SU ($\%Allocated = 50\%$). If, in reality, this project had not already published a paper, the project may have been recommended 90 SU. Because larger recommendations are

reduced more (see reconciliation formula below), this project will be awarded an allocation greater than 45 SU (such that $\%Allocated > 50\%$). Hence, the project with prior publications (and higher expected productivity) has a lower ratio – biasing the results downward.

H.1.1 XRAC Recommendations and Resource Availability

The excerpt below illustrates that reviewers are instructed to make recommendations independent of available XSEDE resources.

> *In all parts of the process, reviewers are asked to consider requests only on their merits. The reviewers are not informed about the total requests on the different systems until the caucus session and only after their individual reviews have been completed. This approach is designed to ensure that reviews are based purely on the merit of the proposal. Providing the allocation request total and available units for each resource at the caucus allows them time to prepare mobility information for requests that need to be or could possibly be moved to alternate systems in an oversubscription situation. The total of the Recommended Allocations is reported to NSF, so that NSF will be aware of the panel's statement of the actual need for resources.*

H.2 Reconciliation formula

After recommendations have been made by the XRAC for each proposal, the Allocations team meets to "reconcile" recommendations with the resources available. In the first ("balancing") phase, the team reallocates resources from over-recommended systems to computing systems which are less over-recommended but are architecturally similar, until the excess demand is balanced equally across platforms. There are generally two categories of architecturally similar resources that XSEDE considered: CPU-oriented systems and GPU-oriented systems.

After the balancing, if there is no longer excess demand for a group of architecturally similar resources, grants are allocated their full recommendation. If there is still excess demand after this balancing occurs, recommended allocations are reduced (the "reconciliation" phase). The amount of the reduction for each project is determined by the following formula:

$$\text{Award} = [(1 - G) \times R_n + (1 - G) \times F \times R_o] \times Size(R)$$

R is the Recommended amount, R_n the share of the Recommended amount supported by NSF-funding (between 0-1) and R_o the share of the Recommended amount not supported by NSF-funding $(1 - R_n)$. F is the funding priority factor which notes how much priority is given to higher NSF shares, and $Size$ is a scaling function that linearly reduces smaller awards less than larger ones.[32] In the formula, G is the global scaling factor, between 0-1, that is solved for in Microsoft Excel to maximize the total of all allocations (Awards) subject to the supply available. The term in brackets is a share less than or equal to one (this can be thought of as a "tax rate"). $Size(R)$ indicates a number of SUs (the recommendation scaled for it's size) that is multiplied by this share to get the final award value.

The process is described in the practices and procedures document as:

> *Following the discussion of all requests, if the aggregate Recommended Allocation is less than the total Available Amount for a system, the Recommended Allocations generally become the Awarded Allocations. When the total Recommended Allocation exceeds the total Available Allocation for a given resource, XSEDE and SP staff use a procedure approved by the NSF to adjust the Recommended Allocations to arrive at the Awarded Allocations. This Reconciliation adjustment is performed by the XSEDE Allocations staff and SP representatives after the XRAC meeting ends and formulaically adjusts all Recommended Allocations to remove oversubscription. Within two weeks of the meeting, the XSEDE allocations staff complete the Reconciliation process, summarize the panel discussions, craft administrative comments or notes to PIs as needed, and send notifications of final awards to all PIs, co-PIs and Allocation Managers of requesting projects, as well as to all affected SPs.*

[32] Sufficiently small recommendations, typically less than one million SU, are not reduced. F takes on a value of either 1 or 0.5 depending on the level of oversubscription. Note that this differs from the XSEDE documentation which lists possible values of F as 0.5 or 2 – this is a typo according to our XSEDE contacts.

I Concepts

Definition

Dimensions AI categorizes and organizes vast amounts of academic research data using machine learning techniques. One of its key features that we utilize in this paper is its ability to extract and identify "concepts" from scientific texts.

Concept Identification: Dimensions AI scans through publications' abstracts and identifies "concepts" associated with each paper. These concepts are meant to represent key themes or topics within a paper. They are derived from the text using machine learning models that recognize important keywords, phrases, and themes commonly associated with specific scientific domains.

Relevance Scoring: Each concept identified by Dimensions is assigned a relevance score, which indicates how central the concept is. While the exact scoring scheme is proprietary, it is a function of the field of study of the paper.

Differences with other methods: The algorithm of Dimensions AI differs from other algorithms such as LDA where the number of topics has to be imputed. This allows us to have variation in breadth. In that sense, Dimensions AI concepts are closer in idea to Hierarchical Dirichlet Process (HDP) that have been used previously in the literature, where the number of topics assigned to a document is not pre-determined.[33]

We present two examples to illustrate the use of concepts.

Paper 1: Defect-engineered graphene chemical sensors with ultrahigh sensitivity

- Abstract: *We report defect-engineered graphene chemical sensors with ultrahigh sensitivity (e.g., 33% improvement in NO2 sensing and 614% improvement in NH3 sensing). A conventional reactive ion etching system was used to introduce the defects in a controlled manner. The sensitivity of graphene-based chemical sensors increased with increasing defect density until the vacancy-dominant region was reached. In addition, the mechanism of gas sensing was systematically investigated via experiments and density functional theory calculations, which indicated that the vacancy defect is a major con-*

[33] See for instance Furman & Teodoridis (2020).

tributing factor to the enhanced sensitivity. This study revealed that defect engineering in graphene has significant potential for fabricating ultra-sensitive graphene chemical sensors.

- https://doi.org/10.1039/C5CP04422G
- Concepts: chemical sensors, conventional reactive ion, density functional theory calculations, ultrahigh sensitivity

Paper 2: Hybrid massively parallel fast sweeping method for static Hamilton–Jacobi equations

- Abstract: *The fast sweeping method is a popular algorithm for solving a variety of static Hamilton–Jacobi equations. Fast sweeping algorithms for parallel computing have been developed, but are severely limited. In this work, we present a multilevel, hybrid parallel algorithm that combines the desirable traits of two distinct parallel methods. The fine and coarse grained components of the algorithm take advantage of heterogeneous computer architecture common in high performance computing facilities. We present the algorithm and demonstrate its effectiveness on a set of example problems including optimal control, dynamic games, and seismic wave propagation. We give results for convergence, parallel scaling, and show state-of-the-art speedup values for the fast sweeping method.*
- https://doi.org/10.1016/j.jcp.2016.06.023
- Concepts: Hamilton-Jacobi equation, algorithm, computer architecture, dynamic game, equations, example problems, fast sweeping method, heterogeneous computer architectures, high performance, hybrid parallel algorithm, optimal control, parallel algorithm, parallel computing, parallel method, parallel scaling, popular algorithms, seismic wave propagation, speedup values, static Hamilton–Jacobi equations, sweeping algorithm, sweeping method, wave propagation