Quantitative Metrics for Generative Justice: Graphing the value of diversity

Métricas cuantitativas de Justicia Generativa: dibujando el valor de la diversidad

Brian Robert Callahan, Charles Hathaway & Mukkai Krishnamoorthy
Rensselaer Polytechnic Institute
callab3@rpi.edu – hathac@rpi.edu – mskmoorthy@gmail.com

ABSTRACT

Scholarship utilizing the Generative Justice framework has focused primarily on qualitative data collection and analysis for its insights. This paper introduces a quantitative data measurement, *contributory diversity*, which can be used to enhance the analysis of ethical dimensions of value production under the Generative Justice lens. It is well known that the identity of contributors—gender, ethnicity, and other categories—is a key issue for social justice in general. Using the example of Open Source Software communities, we note that that typical diversity measures, focusing exclusively on workforce demographics, can fail to fully illuminate issues in value generation. Using Shannon's entropy measure, we offer an alternative metric which combines the traditional assessment of demographics with a measure

of value generation. This mapping allows for previously unacknowledged contributions to be recognized, and can avoid some of the ways in which exclusionary practices are obscured. We offer contributory diversity not as the single optimal metric, but rather as a call for others to begin investigating the possibilities for quantitative measurements of the communities and value flows that are studied using the Generative Justice framework.

KEYWORDS

Demographics; entropy; inclusion; open source software; value.

RESUMEN

Los académicos que trabajan en el marco de la Justicia Generativa se centran principalmente en el análisis cualitativo. Por ello, en este artículo introducimos una medida cuantitativa, la diversidad contributiva, que puede ser utilizada para mejorar el análisis de las dimensiones éticas del valor producido a través de la Justicia Generativa. Es bien sabido que la identidad de las personas que participan de un proyecto—su género, etnia y otras características personales—es un elemento clave para la justicia social. En este sentido, usamos el ejemplo de las comunidades de software de código abierto para destacar que las medidas clásicas de diversidad, que se enfocan exclusivamente en la demografía de la fuerza de trabajo, pueden no ser suficiente para iluminar el problema de la falta de diversidad en los procesos de creación de valor. Usando la medida de entropía de Shannon, desarrollamos una métrica alternativa que combina el tradicional análisis demográfico con la medida de la aportación de valor. Este cruce nos permitirá reconocer contribuciones a la diversidad que normalmente pasan desapercibidas, y evitar la opacidad de las prácticas de exclusión. No obstante, no presentamos la diversidad contributiva como la única medida óptima, sino más bien como una llamada de atención para fomentar la investigación sobre las posibilidades de medir cuantitativamente los flujos de valor que tienen lugar dentro de las comunidades, tal como son estudiados desde el marco de la Justicia Generativa.

PALABRAS CLAVE

Demografía; entropía; inclusión; software de código abierto; valor.

CONTENTS

CONTENIDOS

1. Introduction

Generative Justice seeks to replace the extraction of value, alienated from its generators, with the circulation of value in its unalienated form. It aptly describes, for example, the ways that composting circulates ecological value, worker-owned cooperatives circulate labor value, and online collectives like Wikipedia circulate expressive value. All three are cases in which there is some protection against value extraction by elites from above. But what about injustice from below? How do we know that a bottom-up, collective group has an egalitarian distribution of labor and resources? What might such a measurement look like? Work in the area of Generative Justice has focused primarily on qualitative data; perhaps we are missing out on insights that could be gained if some sort of quantitative data can be produced regarding its internal ethics.

In this paper, we introduce a first attempt at devising what could be quantitative metrics for studying Generative Justice. While we hope this metric is useful, a broader goal is to demonstrate that Generative Justice is something that can in fact be analyzed through quantitative data; that quantitative data can contribute understandings that may not be attainable through solely qualitative methods, and that quantitative metrics can have broad applicability throughout the spectrum of communities studied with a Generative Justice lens. Our measure for this article focuses on a real-world problem: assessment of the diversity of a community that is circulating unalienated value. We have chosen to use the case of Open Source Software (OSS) to apply these metrics: both for OSS's frequent use in the Generative Justice literature, and because diversity is a currently a hot button topic within OSS communities.

In 2010 the GK-12 program at the National Science Foundation provided our Rensselaer Polytechnic Institute with a 5-year grant, *The Triple Helix*, which involved the development of a suite of OSS programs for use in K-12 education and community development (Babbitt et al., 2011; Eglash et al., 2013; Lachney et al., 2016). This allowed us to obtain data in which we could link race and gender identity of each contributor to the code they created. The purpose of this case study was not to test the metric against some independent variable; but merely to explore the process by which the metric could be constructed with real-world data.

Below, we briefly review the definitions of Generative Justice and Open Source Software. We then discuss how diversity is currently measured in OSS projects, and its potential pitfalls.

We contrast this with an examination of metrics measuring value generation. Combining the two, we introduce the concept of contributory diversity. We suggest that this new metric enhances our ability to gain a deeper understanding of how diversity is enacted through the production of value in OSS groups specifically, and that it could potentially be applied to any generative justice system in general.

2. Generative justice

Generative Justice describes an economic framework focused on the bottom-up circulation of unalienated value. The framework is conceived as "orthogonal" to the right/left spectrum of capitalism vs socialism, in that the top-down alienation of value can occur in either regime (extracted to private ownership under capitalism or extracted to state ownership under socialism). Conversely, either system can change to value circulation: it is just as easy to use composted waste, Open Source Software or commons-based media in private institutions as it is in those owned by the state. For the sake of clarity in understanding how we use and understand the term "Generative Justice," we quote its definition:

> The universal right to generate unalienated value and directly participate in its benefits; the rights of value generators to create their own conditions of production; and the rights of communities of value generation to nurture self-sustaining paths for its circulation"

> (Eglash, 2016).

Previous scholarship using the Generative Justice lens includes the Maker Movement (Eglash & Foster, 2014), traditional Japanese wilderness/farm borders (Garvey & Eglash, 2015), STEM education in urban farms (Lyles, 2016), OSS and technical governance (Eglash & Banks, 2014), and children's health education (Bennett et al., 2015). Each of these studies offer cases in which understandings of generative networks involves tracking value flow. But they are all restricted to qualitative, rather than quantitative, data collection and analysis. Generative Justice, we believe, will be enhanced by introducing quantitative metrics that can operate alongside the proliferation of qualitative work.

3. Why open source is relevant

OSS is often cited as an embodiment of Generative Justice: the value created by the software artisan never leaves her hands; it can remain in an unalienated form (code); and it can be freely circulated to other artisans who generate more value. To understand what OSS is, we offer here a brief explanation of how computer programs are designed and implemented. There are two primary pieces to most computer programs: the source code, which represents the program in a human-readable way, and the binary, which is the machine-understood 1's and 0's which comprise the "executable" you run when using an application like Microsoft Word. A good analogy is sheet music: it can be read by artists, and modified by other artists, but once converted into sound it has already been "executed" to produce the useful embodiment so pleasant to our ears. Source code is the sheet music, binaries are the sounds.

Open Source Software is source code which is openly accessible. It is generally accepted that for software to be OSS, it has to meet four requirements: free to use, free to study, free to share, and free to change and share those changes (Free Software Foundation, n.d.). The last requirement is what really makes OSS an excellent example of Generative Justice: musicians can copyright their sheet music, but if they put it in the public domain then others would be free to modify it, reuse bits in other compositions, etc. Indeed, OSS has been so successful that musicians, architects, inventors, designers and other groups have adopted the OSS sharing model for other types of works. Notable in this effort is the Creative Commons (Lessig, 2006) which has much in common with OSS and would serve as another example of Generative Justice. Our metric for examining the diversity of contributions to a software commons would therefore be easily adaptable to these other types of commons-based peer production.

While OSS must be legally available for modification by third parties (rather than just the initial developers), it does not have to be technically easy to modify in order to meet the definition. Much like sheet music, some takes a higher level of knowledge and aptitude to rewrite than others. Without this knowledge, it does not matter whether or not the program is legally OSS, since it is effectively closed source to the person who cannot understand it. In such cases it could be said that the projects are "OSS in name only," that is, they are not easily modified by others, and ultimately do not allow for the benefits of Generative Justice to cycle through the system.

In example of projects that may be OSS in name only are applications (apps) available through the Apple iOS App Store that are licensed under an OSS license. First, the odds of

coming across an OSS app in the Apple store are miniscule: as of April 2016 there were only 110 OSS apps out of the over 2 million apps available (Open Source iOS Apps —Real iOS Source Code Examples, 2016). Second, even if a user beat the odds and found one, the App Store does not link the app to its source code, so in order to benefit from the OSS status she would have to search the web to find the app on a separate site that does offer that link. Although Open Source code is free, modifying that code for use on the Apple OS is not: you have to obtain a proprietary development environment from Apple. Finally, Apple does not permit you to freely transfer apps you create to more than 5 additional machines: after that you have to go through Apple's vetting process and offer it from their store (even if it's free). As best there is a considerable delay during the review process, and ultimately the ability for any of these modified apps to be added back to the App Store is not certain.

In other words, some OSS projects better match the egalitarian ideals of Generative Justice than others. Metrics which can help guide us towards Generative Justice will be sensitive to exclusionary practices, and help identify those systems in which participation in the creation and benefits of value creation and circulation are spread as broadly and democratically as possible.

4. Diversity issues in OSS

Doing tech in the 21st Century requires more than just a computer and a desire to learn. As we come to rely increasingly on distributed networks of loosely affiliated programmers making software together—the organizational method of OSS—questions about what it means to be a part of these networks, who can be a part of these networks, and the ability to accurately measure and portray their demographics are increasingly tied to issues of social justice. Despite the history of computing as primarily white and male (Levy, 1984), women and underserved ethnic groups are now striving for equity. However, current methods for understanding the workings of computer programmers engaged in producing OSS have left out many of the human questions of technical production. Qualitative approaches such as ethnography, though excellent for interpreting and understanding much of the interpersonal dynamics in these large and complex systems, could be improved if they were complimented by quantitative metrics using demographic data. Even when academics do pursue the inner workings of particular OSS networks, they can miss such important factors as diversity: a perusal of Gabriella Coleman's (2013) book *Coding Freedom*, which gives an ethnographic

account of the world's largest OSS project, contains in its index neither a category for "diversity", nor "gender", nor "race", nor "women." "Debian, female representation in" (Debian being the name of the OSS project Coleman studied) received four total pages, spread throughout the book. It seems necessary, then, that we ought to begin to give attention to quantitative methods for diversity metrics, if we are in fact desiring to make assessments regarding diversity within OSS.

The ability to easily obtain a grasp of the diversity within OSS, to say nothing to the tech industry at large, has had its share of difficulties. Indeed, this was the observation by one of the authors, who in a previous study realized that he could only at best perform a very weak approximation of diversity of the OSS project he was studying, via speaker lists from the community's public conferences (Callahan, 2016). Fortunately, even with these difficulties, there has been academic work into uncovering these insights. In one study, a pair of surveys administered approximately a decade apart, there were positive trends in the increasing diversity of Open Source as a whole, at least along a binary axis of gender. The original survey, administered in 2002 found that less than 2% of respondents identified as non-male (Ghosh et al., 2002); its follow-up survey administered in 2013 discovered that those who identified as non-male jumped to approximately 14% (Arjona-Reina et al., 2014).

Some of the change is due to collaborations between the computing community and social justice efforts. Several underserved-only OSS groups have formed in the last several years: perhaps most notably Outreachy[1] and PyLadies, the latter of which has helped the Python OSS community increase their female speaker ratios and general attendance ratios at the Python OSS conference PyCon from approximately 11% to over 30% in just a few short years (Root, 2014). Furthermore, those who are at the margins are beginning to self-organize independent of existing communities in order to gain enough power and momentum to fight for equity on their own terms. One such example is the collective of transgender women that have organized to fight for the normalcy of transgender hackers in OSS (Callahan and Lemmer, 2016).

In addition to diversity as an embodiment of social justice, it also has utilitarian features: there is a significant literature that shows positive correlations between diversity and innovation (for example: Dell'Era and Verganti, 2010; Florida and Gates, 2003; Gassmann, 2001; Van der Vegt and Janssen, 2003; Østergaard et al., 2009). Additionally, the circulation of unalienated value of any system, not just those of OSS, depends on connectivity.

[1] <https://www.gnome.org/outreachy/>

Connectivity can be enhanced by diversity, a phenomenon which is evident for both human and non-human agents, as we see for example in marine ecosystems (Mora et al., 2016).

5. How do we typically measure diversity?

Diversity as usually understood in the technical field is based on comparisons between the number of people in a given identity category in the workforce under investigation, and the percentage of that identity category in the general population. For example, if women are 50% of the population, they should be 50% of the workforce.

Without negating the validity of this measurement, we offer four critiques suggesting that an alternative would also have merit. Figure 1 shows data from our case study, in which we see the percentage of participants in each ethnic group. But there is no information here on the activity of each individual: we don't know which participants took on which roles; if some dominated certain resources or tasks, and if others felt they needed to hold back. That is, how does the value generated by each group compare to its demographic presence?

Figure 1. Percent of 3Helix Population vs Ethnicity
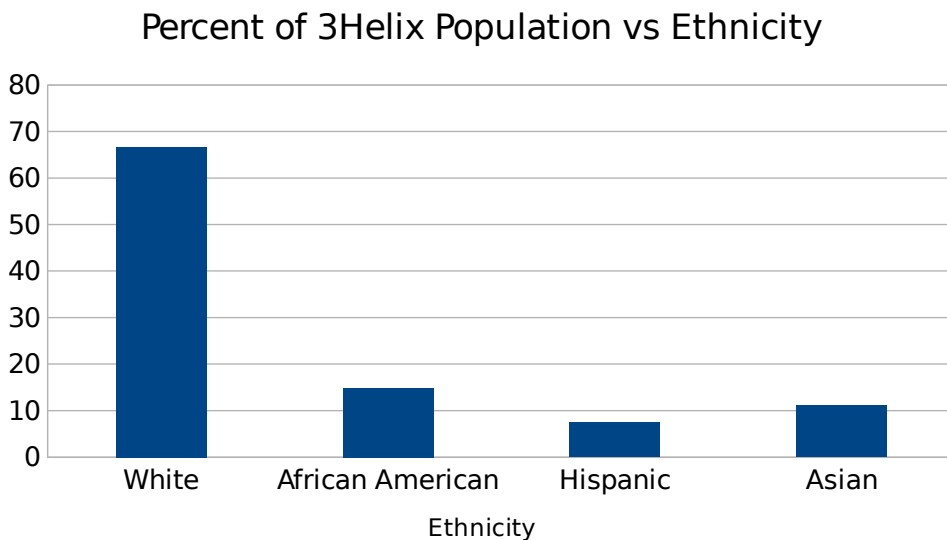
Source: Data collected for study.

Second, even if we do pay attention to the demographics of particular roles, recent scholarship has questioned its relationship to social justice issues. Assereckova (2016) notes that Latvia is lauded for having women occupy 36% of the top business positions such as CEO. But the total number of CEOs is only a tiny percentage of the female workforce, and in Latvia women in the working class have a large gender wage gap. Here is a case where attention to production roles is masking inequality, not aiding in its exposure. Indeed, one could imagine someone with a fabricated title such "Senior Test Engineer in Charge of Software Configuration" doing nothing but writing configuration settings that later get replaced over and over again, and generally contributing nothing to the project. The Latvian case as well as the problem of false titles point to the necessity for quantitative metrics to be undergirded by an intersectional approach (Crenshaw, 1989) that takes into account how labor is extracted, alienated or represented.

A third issue is illustrated by the case of black medical students. In contrast to the debates over affirmative action, whose critics maintain that resources are wasted if efforts are helping to increase the diversity in highly competitive admissions such as medical school, the American Medical Student Association points out that encouraging diversity in the medical professions succeeds in producing doctors that go on to practice in low-income communities (AMSA n.d.).

The lack of medical access in these communities create a variety of social ills (poor school attendance, underemployment, etc.). In other words, the value generated and circulated in such cases is poorly tracked by merely focusing on job titles, income levels, or other typical demographic approaches to diversity.

Finally, there are concerns over gender bias in the acceptance of code in OSS. A recent article which uses GitHub pull requests as a measure of the diversity within an OSS project found that, while women are more likely than men to have their code contributions accepted into a project, that likelihood is only true so long as the gender of the woman is not known. Once she has been identified as a woman, her code contributions are far less likely to be accepted and far more likely than a man's to be rejected (Terrel et al., 2016).

In all four cases, the issue is in keeping track of both demographic diversity as usually conceived, and simultaneously tracking the value that is generated and circulated. Below we propose a metric which combines those two parameters.

6. Contributory diversity: towards a new quantitative metric

To account for these issues, one could imagine measuring the *contribution* each person puts into a project. This does not have to translate directly to revenue—anything that contributes to the generation of value could be considered—and while it might be specific to the domain, it would still provide a useful tool to garner insights. In OSS the number of lines of code, number of pull requests, or other parameters reflecting the amount of contributions made to the software offer good representations for this value generation. Because our case study had contributions spanning across a broad array of software and hardware projects, we tracked blog posts from project participants as a stand-in for contributions to software development. Literature on software development suggests that communications such as blog posts and email can be reliable indicators of code contributions (Pagano and Maalej, 2011; Valverde et al., 2006). We chose the 3Helix program because this provided us with the gender and ethnicity for each person posting to the blog.

Recall that Figure 1 showed the number of project participants in each ethnic category. Figure 2 presents the number of blog posts for each ethnic category. This does not take into account the number of people in each ethnic category. In an ideal, egalitarian project, each ethnic group would contribute the percentage of value proportionate to their population.

Combining Figure 1 and Figure 2 produces Figure 3, the number of contributions in each ethnic category divided by percentage in the workforce population. Figure 3 visually represents the internal contributory diversity metric.

Finally, we can create an external contribution metric by normalizing the identity category percentages to society at large rather than just internally. This is shown in Figure 4.

Now that we have a visualization for contributory diversity, we can develop a metric for it. Here we use entropy. When entropy is maximized, distributions are as random as possible, thus reflecting only the percentage of each identity category in the population. This is why entropy is used in ecological studies of diversity: a highly diverse ecosystem would have approximately equal numbers of each species and this would have the maximum possible entropy.

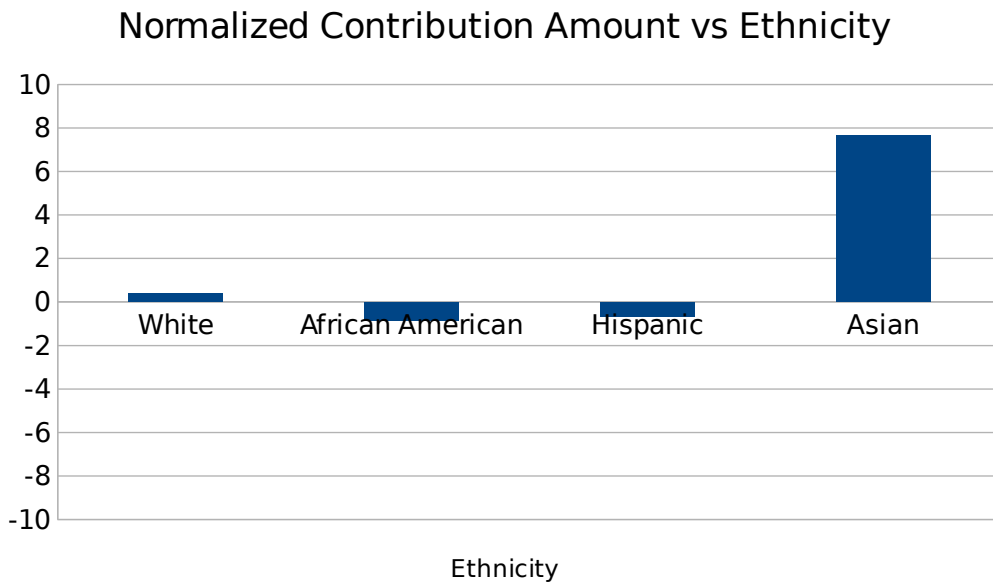FIGURE 2. TOTAL CONTRIBUTIONS VS ETHNICITY
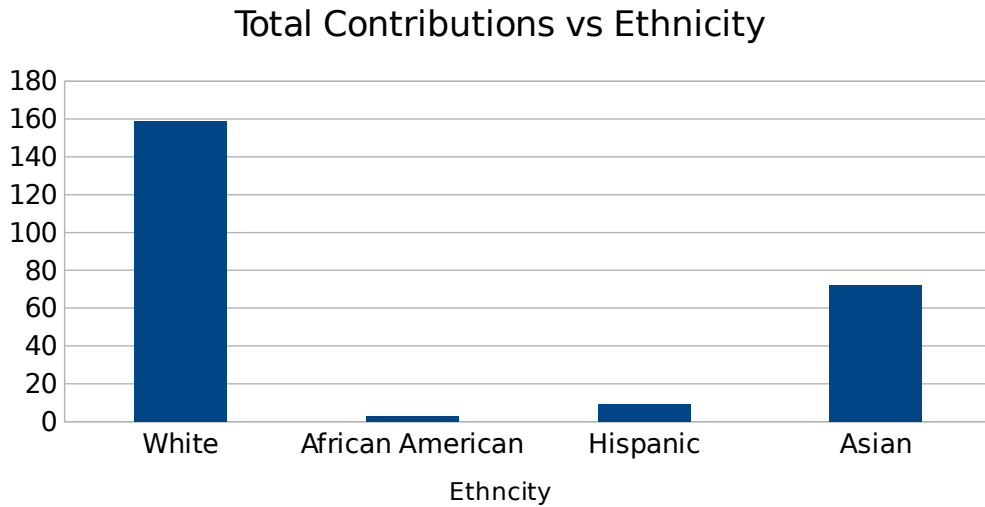
Source: Data collected for study.

FIGURE 3. PERCENTAGE OF CONTRIBUTIONS VS ETHNICITY
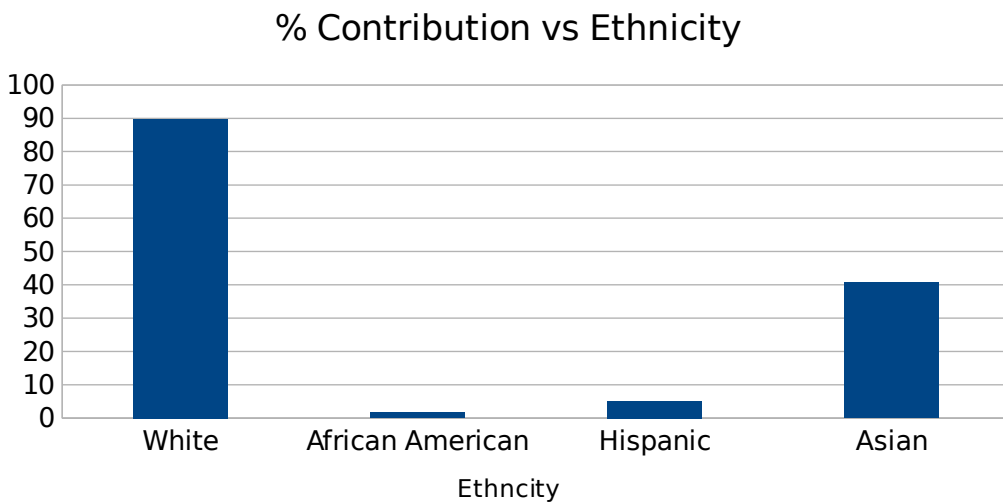
Source: Data collected for study.

Figure 4. Normalized Contribution Amount vs Ethnicity

Source: Data collected for study.

We build upon previous attempts to measure diversity in other fields, for example biology and information theory. One of the most well-known metrics is Shannon Entropy (Shannon, 1948). It is used in biology to measure diversity within an ecosystem, within information theory to quantify the number of bits in a given string of data, and we apply it here to diversity within a workforce.

An example of Shannon Entropy would be representing how difficult it is to predict the color of the next marble drawn from a bag. The probability of any color is equivalent to the prevalence of the color (i.e. the percentage of the population with that color). If four colors are equally distributed, then the entropy is at its maximum, because each one is equally probable. Shannon calculated the entropy by taking the base-2 log of the reciprocal of the probability:

So the entropy of the group as a whole—the average entropy per category—is the weighted sum of those values (that is, weighted by multiplying by that probability):

In the case of 4 marble colors, the maximum diversity would be -4*(0.25*ln(0.25)) = 2. Moving from marbles to humans, is the probability that a given ethnicity would be selected at random (which is simply the percentage of that ethnicity in the population). For our data-set, these are the results:

- Shannon Index for population ethnicity: 0.9901331507

- Shannon Index for contribution ethnicity: 0.6828124194

- Shannon Index for percent of contributions by ethnic representation: 0.2210632337

To help interpret these results, a value of 0 would mean absolutely no diversity. For example, if the population was entirely white, we would be guaranteed that any coder selected at random would be white, so the entropy would be 0. If the population was distributed equally in all categories, we would have the maximum entropy. With the population of the 3Helix project, there are four possible identity categories, so the maximum possible would be 2. The data shows a moderate amount of diversity (represented by the approximately 0.99 population ethnicity Shannon Index). However, the population ethnicity metric does not map perfectly onto the contributions made by each ethnicity, as represented by the approximately 0.68 contribution ethnicity metric. This is a direct result of the underrepresentation of the African American and Hispanic contributions in the 3Helix project.

The last result, metric for percent of contributions by ethnic representation, is even lower because it is compounded by the other two results: the slightly disproportional population and the highly disproportional contributory diversity. This final result indicates that the contributions in the 3Helix project did not correlate to what we would expect in an ideal project. That is to say, the 3Helix project had low contributory diversity.

7. Conclusion: the future of quantitative analysis in generative justice

Quantitative measurements that can be used to study Generative Justice is an area of research that has heretofore gone without study. And certainly contributory diversity is not the only quantitative metric that can be devised for analysis with a Generative Justice framework. We hope that contributory diversity, besides being useful in its own right, will encourage others to produce quantitative metrics of their own. Shifting Generative Justice towards analytic efforts that incorporate both qualitative and quantitative data collection and analysis can only strengthen the framework and its insights.

Contributory diversity brings the issues of value generation into conversation with the demographic measures that usually dominate such studies. The diversity of a particular workforce can be viewed through the lens of actual contributions, and not merely broad identity categories or job titles. Given the efforts to have industry meaningfully increase diversity in the tech sector, and the social justice issues which diversity can embody, a statistic that can help ensure that the members of this diverse workforce are equals in generating and circulating value, a core insight of the Generative Justice framework, becomes critically necessary.

In future work, we hope that a general interest in quantitative metrics for Generative Justice could be applied to institutions, pathways, value chains and other aspects that reach beyond the level of individuals and toward a more comprehensive look at the activities of networks and groups that are engaged in this bottom-up, self-organized unalienated value flow.

8. References

AMSA. (n.d.). *Enriching Medicine Through Diversity.* Retrieved from
<http://www.amsa.org/about/mission-aspirations/diversity/>

ARJONA-REINA, L., ROBLES G., & DUEÑAS, S. (2014). *The FLOSS2013 Free/Libre/Open Source Survey.* Retrieved from
<http://floss2013.libresoft.es/results.en.html>

ASSERECKOVA, M. (2016). *Women and the Economy: Why Women Becoming CEOs Might Not be the Best Benchmark for Feminism*. European Young Feminists: A Platform for Young Feminists in Europe. Retrieved from
<http://www.europeanyoungfeminists.eu/2016/01/31/women-and-the-economy-why-women-becoming-ceos-might-not-be-the-best-benchmark-for-feminism/>

BABBITT, W., LYLES, D., & EGLASH, R. (2011). From Ethnomathematics to Ethnocomputing: Indigenous Algorithms in Traditional Context and Contemporary Simulation. In S. Mukhopadhyay and W.-M. Roth (Eds.), *Alternative Forms of Knowing (in) Mathematics: Celebrations of Diversity of Mathematical Practices* (pp. 205-220). Rotterdam: Sense Publishers.

BENNETT, A., CASSIM, F., VAN DER MERWE, M., VAN ZIJL, K., RIBBENS, M. (2015). Designing South African Children's Health Education as Generative Play. *Proceedings of the International Association of Societies of Design Research, Queensland University of Technology*, November 2015

CALLAHAN, B. R. (2016). SocialBSD: A Review of Diversity and Inclusivity Initiatives in the *BSD Community and Imagining Future Pathways. Proceedings of AsiaBSDCon 2016 (pp. 167-176).

CALLAHAN, B. R. & LEMMER, L. (2016). Trans Code: Free Software as Model & Critique of Diversity by Transgender Hackers. Presented at "LibrePlanet 2016," MIT, March 19-20, 2016.

COLEMAN, E. (2013). *Coding Freedom: The Ethics and Aesthetics of Hacking*. Princeton: Princeton University Press.

CRENSHAW, K. (1989). Demarginalizing the Intersections of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics. *The University of Chicago Legal Forum, 140*, 139-167.

DELL'ERA, C. & VERGANTI, R. (2010). Collaborative Strategies in Design-Intensive Industries: Knowledge Diversity and Innovation. *Long Range Planning, 43*(1), 123-141.

EGLASH, R. (2016). An Introduction to Generative Justice. *Teknokultura,* 13(2).

EGLASH, R. (2014). Generative Justice: The Revolutions Will Be Self-Organized. *Tikkun,* Retrived from
<http://www.tikkun.org/nextgen/generative-justice-the-revolution-will-be-self-organized>

EGLASH, R & BANKS, D. (2014). Recursive Depth in Generative Spaces: Democratization in Three Dimension of Technosocial Self-Organization. *The Information Society, 30*(2), 106-115.

EGLASH, R. & FOSTER, E. On the Politics of Generative Justice: African Traditions and Maker Communities. Presented at "What do Science, Technology, and Innovation Mean from Africa?" MIT, November 13-15, 2014.

EGLASH, R., GILBERT, J.E. & FOSTER, E. (2013). Toward Culturally Responsive Computing Education. *Communications of the ACM, 56*(7), 33-36.

EGLASH, R., BENNETT, A., O'DONNELL, C., JENKINS, S. & CINTORINO, M. (2006). Culturally Situated Design Tools: Ethnocomputing from Field Site to Classroom. *American Anthropologist, 108*(2), 347-362.

FLORIDA, R. & GATES, G. (2003). Technology and Tolerance: The Importance of Diversity to High-Technology Growth. In Terry Nichols Clark (Ed.) *The City as an Entertainment Machine (Research in Urban Policy, Volume 9)* (pp. 199-219). Boston: Elsevier.

FREE SOFTWARE FOUNDATION. (n.d.) "What Is Free Software? - GNU Project - Free Software Foundation." Retrieved from
<http://www.gnu.org/philosophy/free-sw.html>

GARVEY, C. & EGLASH, R. (2015). Satoyama as Generative Justice: Humanity, Hybridity, and Biodiversity in a Horizontal Social Ecology. In E. Marhöfer and M. Lylov (Eds.), *Shape Shifting* (pp. 58-68). Berlin: Archive Books.

GASSMANN, O. (2003). Multicultural Teams: Increasing Creativity and Innovation by Diversity. *Creativity and Innovation Management, 10*(2), 88-95.

GHOSH, R. A., GLOTT, R., KRIEGER, B., & ROBLES, G. (2002). Free/Libre and Open Source Software: Survey and Study, Part 4: Survey of Developers. International Institute of Infonomics, University of Maastricht. Retrieved from
<http://flossproject.org/report/FLOSS_Final4.pdf>

LACHNEY, M., BABBITT, W., & EGLASH, R. (2016). Software Design in the "Construction Genre" of Learning Technology: Content Aware Versus Content Agnostic. *Computational Culture, 5*. Retrieved from
<http://computationalculture.net/article/software-design-in-the-construction-genre-of-learning-technology-content-aware-versus-content-agnostic>

LESSIG, L. (2006). *Code: Version 2.0*. New York: Basic Books.

LEVY, S. (1984). *Hackers: Heroes of the Computer Revolution*. Garden City, NY: Anchor Press/Doubleday.

LYLES, D. (2016). *Generative Contexts*. Unpublished doctoral dissertation, RPI, Troy, NY.

MORA, C., DANOVARO, R., & LOREAU, M. (2014). Alternative Hypotheses to Explain Why Biodiversity-Ecosystem Functioning Relationships are Concave-Up in Some Natural Ecosystems but Concave-Down in Manipulative Experiments. *Scientific Reports 4*, Article Number 5427. Retrived from
<http://www.nature.com/articles/srep05427>

"OPEN SOURCE IOS APPS — REAL IOS SOURCE CODE EXAMPLES." (2016). Retrieved from
<https://maniacdev.com/2010/06/35-open-source-iphone-app-store-apps-updated-with-10-new-apps>

ØSTERGAARD, C. R., TIMMERMANS, B., & KRISTINSSON, K. (2011). Does a Different View Create Something New? The Effect of Employee Diversity on Innovation. *Research Policy, 40*(3), 500-509.

PAGANO, D. & MAALEJ, W. (2011). How Do Developers Blog? An Exploratory Study. Proceedings of the 8th Working Conference on Mining Software Repositories (pp. 123-132).

ROBLES, G. & GONZALEZ-BARAHONA, J. M. (2006). Contributor Turnover in Libre Software Projects. In E. Damiani, B. Fitzgerald, W. Scacchi, M. Scotto, and G. Succi (Eds.). *IFIP International Federation for Information Processing, Volume 203, Open Source Systems* (pp. 273-286). Boston: Springer.

ROOT, L. (2014). Fundraising for PyLadies for PyCon 2015. PyLadies. Retrieved from
<http://www.pyladies.com/blog/fundraising-for-pycon-2015/>

SHANNON, C. E. (1948). A Mathematical Theory of Communication. *The Bell System Technical Journal, 27*, 379-423 and 623-656.

Terrell, J., Kofink, A., Middleton, J., Rainear, C., Murphy-Hill, E., Parnin, C. & Stallings, J. (2016). Gender bias in open source: Pull request acceptance of women versus men. PeerJ PrePrints. 4:e1733v2
<https://doi.org/10.7287/peerj.preprints.1733v2>.

Valverde, S., Theraulaz, G., Gautrais, J., Fourcassie, V., & Sole, R. V. (2006). Emergent Behavior in Agent Networks: Self-Organization in Wasp and Open Source Communities. Retrived from
<https://arxiv.org/ftp/physics/papers/0602/0602003.pdf>

Van der Vegt, G. S. & Janssen, O. (2003). Joint Impact of Interdependence and Group Diversity on Innovation. *Journal of Management, 29*(5), 729-751.