

Using AI To Help Job Seekers

Lessons from SkillUp AI

JULY 2025

Career coaching provides value to job seekers, but it is labor-intensive, requires knowledge of many career options and localities, and is often difficult for job seekers to navigate.¹ These limitations leave many job seekers—particularly those with fewer resources to find and engage with coaches—at a disadvantage in the labor market. The recent rise of generative artificial intelligence (AI) may present an opportunity to provide low-cost, personalized coaching to job seekers on a large scale. Some early research into using AI to assist with career coaching has found promising results.²

[SkillUp Coalition](#) is a nonprofit organization dedicated to helping individuals skilled through alternative routes—meaning people who have skills obtained outside of college—secure employment with high opportunities for career advancement that can be achieved regardless of a college degree. Through its online platform and step-by-step career guidance options, SkillUp has connected more than 3.5 million workers to career support, training, and job opportunities. In collaboration with its national and local partners, SkillUp curates career paths, training programs, coaching resources, and job listings that lead to living-wage, high-opportunity careers. SkillUp’s website hosts this catalog of materials, allowing job seekers to browse opportunities.

*Bennett Otten (MDRC)
Kelsey Schaberg (MDRC)
with
Jason Cheng (SkillUp Coalition)*



While SkillUp's current platform provides users with information, it does not yet provide all the personalized support and information job seekers could benefit from in their careers. Generative AI may be able to offer such personalized guidance to job seekers, which could lead to higher engagement with the varying populations using SkillUp, and help them to find new opportunities. An AI-enabled resource on SkillUp's platform could help its users navigate its catalog in a way that connects with them, allowing SkillUp to reach more people and meet new demands more easily. SkillUp therefore codesigned and tested an AI chatbot to present opportunities to users from its catalog. This brief describes that test and its results.

SkillUp hoped that the chatbot would accelerate the career-exploration process and reduce the number of searches needed to find opportunities.

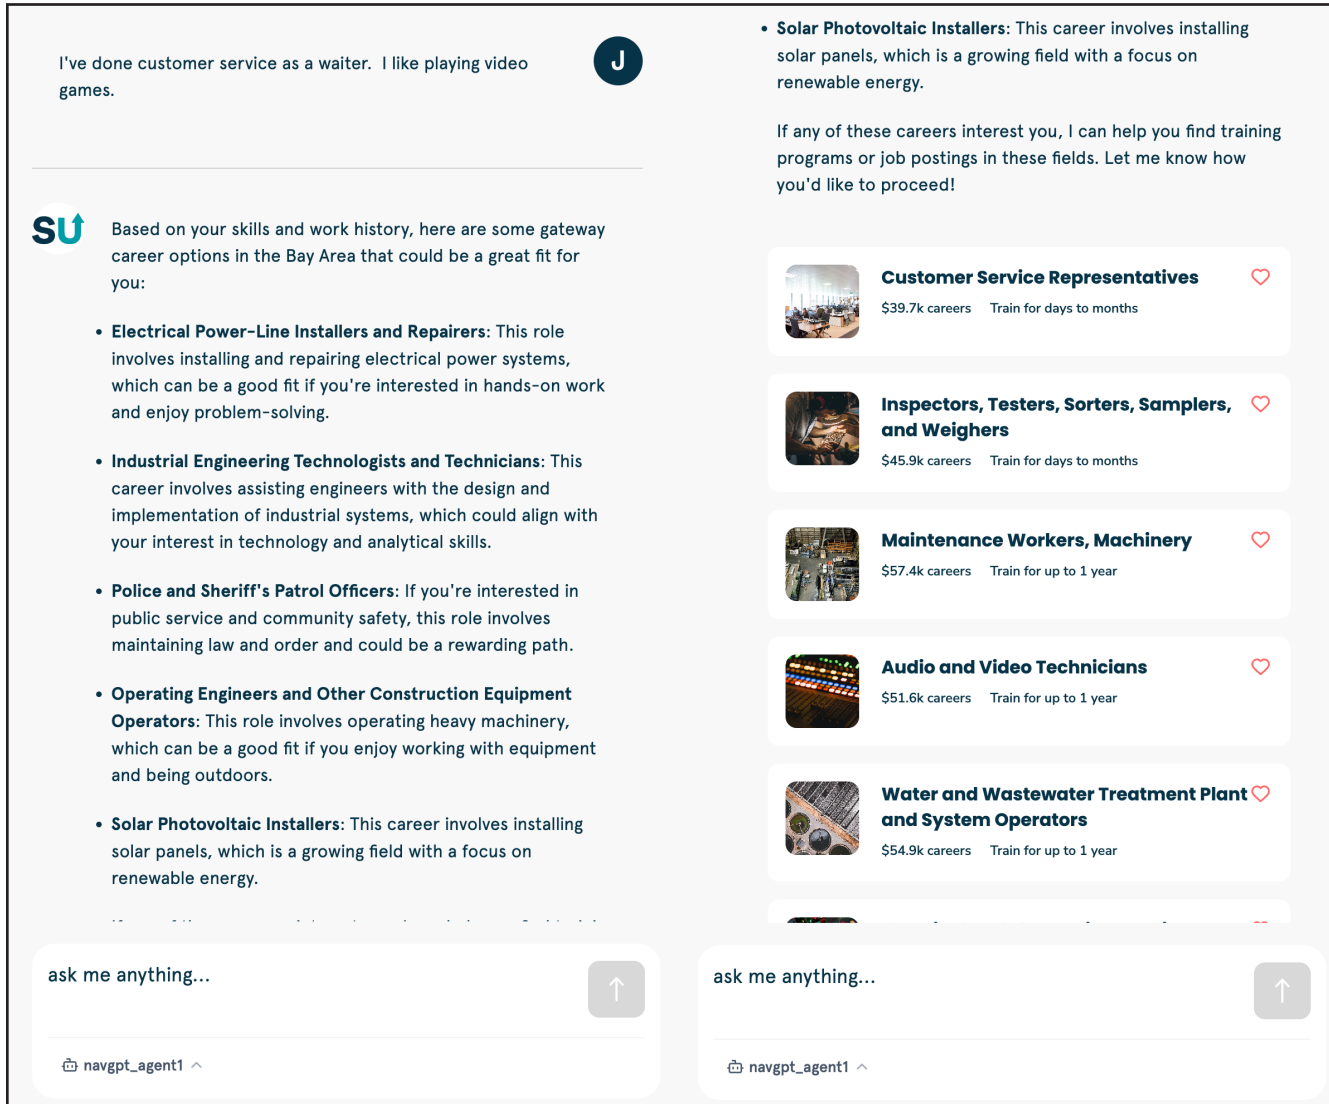
Developing SkillUp AI, a Career Coaching Chatbot

An early adopter of AI technology, SkillUp set out to draw on the power of existing large language models (LLMs) — machine learning models that can generate text — to create a chatbot that could provide job seekers with personalized recommendations. This benefit to users would be shown by an increase in the number of users following through to additional stages in their job searches, for example the number users seeing relevant items in SkillUp's catalog and clicking to external websites to enroll in training or apply for a job, compared with the number following through from the main SkillUp platform (called "SkillUp Explore"). SkillUp also hoped that the chatbot would accelerate the career-exploration process and reduce the number of searches needed to find opportunities.

To create the chatbot, known as SkillUp AI, SkillUp and its technology partner WhereWeGo secured funding from several foundations in partnership with Brighthouse, which provided the AI infrastructure, and AdeptID, which provided a recommendation engine that matches user-provided skills and interests to the curated selections in the SkillUp catalog. The chatbot, based on a version of OpenAI's ChatGPT, converses with users about their skills, job history, and other relevant topics before calling the recommendation engine to present users with specific careers, training programs, or job listings from SkillUp's catalog based on that input. SkillUp AI's chatbot views SkillUp's curated listings and uses them as the primary basis for responses.³ Figure 1 shows an example conversation with SkillUp AI.

SkillUp pilot tested its chatbot with 500 users in November 2024. The pilot users provided real-world testing of the tool and were sent a short survey about SkillUp. They set a baseline for how users were likely to use SkillUp. SkillUp used this information to make improvements to the tool. These improvements included bug fixes — for example, fixing an error that prevented the chatbot from providing responses if a user did not provide a work history and letting users switch the locations for which they received recommendations. SkillUp also worked with its partners to improve the recommendation engine, which favored certain recommendations regardless of a user's inputs. The pilot test also revealed implementation challenges for deploying the chatbot. Difficulties in connecting the chatbot infrastructure to the recommendation engine sometimes caused slow or imperfect responses. SkillUp also faced difficulties coordinating changes across multiple organizations with different work schedules and priorities.

Figure 1. Example of a SkillUp AI Conversation



Testing the Effectiveness of SkillUp AI

SkillUp then began an “A/B test” of SkillUp AI, with analysis performed by MDRC, to see whether the chatbot led to better engagement and career-related outcomes for users than the main SkillUp platform. A/B tests compare two different versions of something – in this case SkillUp AI and SkillUp Explore – to see which version leads to better results. In this experiment, MDRC examined whether users received and followed through on recommendations for career paths, training, and jobs.

For the test, job seekers were directed to a sign-up page on SkillUp by an ad. Once they created an account, they were randomly assigned to one of two groups. The first group (called the “SkillUp Explore group”) had access to SkillUp’s website as it existed before the development of SkillUp AI. Users in this group were able to view the opportunities in SkillUp’s catalog through browsing and searches, along with optional guidance that introduced them to different features of the website. When users saw opportunities in which they had interest, they could save those opportunities to their profiles or view additional details about the opportunities, including links to external sites where they could view additional information or apply. The second group (called the “SkillUp AI group”), had access to a site containing only SkillUp AI, where they were able to send and receive messages with the chatbot but were unable to browse SkillUp’s catalog. Instead, SkillUp AI’s recommendations were shown as cards with links to opportunities, which users could save or navigate to if they had interest.

Between February 15 and March 7, 2025, 5,036 job seekers were randomly assigned.⁴ The sample is majority female, racially diverse, and skews young, with the majority of the sample under 35 years old. A majority of the sample had a high school diploma or less, and just under 60 percent had at least five years of work experience. This sample largely matches the demographics of SkillUp’s user base in recent years, with the only major difference being that the A/B test sample was slightly younger. Random assignment was successful, and there were no statistically significant differences in collected demographics between the two groups.

Findings

The findings presented here draw on internal data collected by SkillUp and shared with MDRC.

- **Compared with the SkillUp Explore group, the engagement for the SkillUp AI group was significantly lower. There is also little evidence of SkillUp AI leading users to training and job opportunities more directly.**

Participants in the SkillUp AI group were 40 percentage points less likely to view a catalog item and 18 percentage points less likely to save a catalog item (Table 1) than participants in the SkillUp Explore group. SkillUp AI users were also less likely to pursue opportunities, with only 4 percent of users clicking to external websites to apply for training and jobs compared with 30 percent of SkillUp Explore users.

- **Over 50 percent of SkillUp AI users did not advance far enough in their conversations with the chatbot to receive a recommendation, meaning they did not view opportunities within SkillUp.**

To receive recommendations, users in the SkillUp AI group were prompted by the chatbot for information on their work history, location, and skills. Many users declined to provide this information. Around 45 percent of users in the SkillUp AI group sent two or fewer messages. This pattern suggests a hesitancy to use chatbots, as has also been observed in other research.⁵

Table 1. Impacts of the A/B Test

Outcome	SkillUp Explore	SkillUp AI	Impact
Used SkillUp after sign-up (%)	94.9	84.3	-10.6 ***
Ever performed an action (%)			
Viewed a catalog item	59.5	17.9	-41.6 ***
Saved a catalog item	24.6	5.7	-18.9 ***
Left SkillUp to apply	29.6	3.9	-25.7 ***
Average number of actions performed			
Catalog item views	2.1	0.4	-1.7 ***
Catalog item saves	1.4	0.1	-1.3 ***
Departures from SkillUp to apply	0.5	0.0	-0.4 ***
Average time on SkillUp (minutes)	8.1	3.7	-4.4 ***
Sample size (total = 5,036)	2,494	2,542	

SOURCE: MDRC calculations using SkillUp tracking data.

NOTE: A two-tailed t-test was applied to estimated differences. The statistical significance is indicated as follows: *** = 1 percent; ** = 5 percent; * = 10 percent.

- **The differences in viewing, saving, and acting on opportunities are much smaller when comparing SkillUp AI users who received a recommendation and SkillUp Explorer users.**

While this comparison, shown in Table 2, is not experimental, the increase in engagement relative to the full group shows that some of the differences between the study groups are attributable to participants' unwillingness to use an AI chatbot, rather than the chatbot's performance in presenting people with relevant recommendations. However, even within this subgroup who received recommendations, engagement in the SkillUp AI group remains much lower than that in the SkillUp Explore group.

Table 2. Outcomes Among SkillUp Explore Users and SkillUp AI Users Who Received Recommendations

Outcome	SkillUp Explore	SkillUp AI + Recommendation	Difference
Ever viewed a catalog item (%)	59.5	38.2	-21.3
Career	25.5	26.8	1.3
Training	24.4	11.9	-12.5
Job	37.7	10.4	-27.3
Ever saved a catalog item (%)	24.6	12.3	-12.3
Career	15.2	9.0	-6.2
Training	10.0	3.5	-6.4
Job	8.3	1.4	-6.8
Ever left SkillUp to view an opportunity (%)	29.6	7.4	-22.2
Training	12.3	4.6	-7.8
Job	19.6	3.0	-16.6
Average time on SkillUp (minutes)	8.1	6.1	-2.0
Sample size	2,494	1,184	

SOURCE: MDRC calculations using SkillUp tracking data.

Understanding the Results

In What Ways Did SkillUp AI Show Promise?

- **SkillUp AI was able to provide some users with personalized recommendations and had the most success recommending relevant career paths.**

Some aspects of SkillUp AI showed promise. While the participants in the SkillUp AI group who received a recommendation were less likely to click a link to external training or a job posting, they viewed and saved career catalog items at a rate similar to that of the SkillUp Explore group. This similarity may show that while SkillUp AI was not as effective at producing relevant training and job recommendations, it demonstrated potential at recommending relevant career options.

Records of the conversations with SkillUp AI show some examples where the chatbot functioned very well. For example, one participant described having experience in manufacturing, and the recommendation engine then returned listings in the area that would allow that user to enter similar fields with more opportunities. The participant then asked the chatbot for more specific opportuni-

ties in certain career types, viewing multiple recommended options before clicking on an external website to view an application for training. This interaction shows the potential of SkillUp AI to make connections organically, while also allowing the user to guide the chatbot in a desired direction.

Why Might Engagement with SkillUp AI Have Been Low?

- **SkillUp AI users had to send several messages before they received or could view opportunities.**

The amount of time it took to receive a recommendation may have led many users to stop interacting with the chatbot, and probably continued to affect the experience of those who received recommendations. Some users also experienced delays between sending a message and receiving a recommendation, which may have led to them to exit the chat before receiving a recommendation.⁶ Some SkillUp AI users were not shown the type of opportunities they wanted to explore. Compared with SkillUp Explore, far more of the recommendations for SkillUp AI users were for general career paths, rather than specific training opportunities or jobs.

Participants in the SkillUp Explore group could click on different tabs to see only the opportunities that matched their interest. On SkillUp AI, it was less clear to participants what kinds of opportunities were available, or that they could request different kinds of opportunities to be shown. This lack of clarity probably contributed to the decrease in the number of external job and training website visits and also may have caused frustration for SkillUp AI users.

- **Some SkillUp AI users probably received career, training, and job recommendations that did not match their interests, leading them to stop interacting with the chatbot.**

SkillUp Explore participants could view any opportunities that were available in their location. SkillUp AI participants received a few recommendations at a time, based on their responses to the chatbot. Sometimes, these recommendations matched the skills and interests gleaned from the participants' responses to earlier questions. Other times, the recommendation engine did not have much information to go on (since the user had not provided much), leading to recommendations in which the user had no interest.

Occasionally, participants asked for opportunities in specific sectors or occupations that either were not in SkillUp's catalog as gateway careers or were unavailable in the user's geographic location. In some cases, these issues may have been solvable had the user kept prompting the chatbot, but entering these additional prompts would have taken real effort. It appears that on many occasions when users did not receive helpful recommendations initially, they often stopped interacting with the chatbot.

Lessons from the SkillUp AI Evaluation

There are plenty of lessons to learn from this experiment.

- **AI chatbots have some disadvantages compared with standard websites.**

Most companies that operate primarily through websites have expertise in creating webpages that are rich in information and easily digestible. Modern web design is well understood and enables users to easily browse and filter information with just a few clicks. While entering text into a chatbot can seem more intuitive, it may instead add friction not present in a well-designed website. At present, chatbots should not be viewed as a replacement for existing websites, though they may have uses as a supplement to current features.

- **Using a chatbot effectively may be harder than many people initially believed.**

When developing a chatbot, it is easy to assume that users will know what to ask. However, prompt engineering — knowing how to write a message in order to get the desired information back — is a skill that many people do not currently have. Typing a precise question into a chatbot is not intuitive for everyone, especially when it is not clear what to ask. SkillUp believes it underestimated how hard this task would be for users.

- **AI chatbots may be better suited to guided exploration, where users are less sure of their direction and want assistance.**

One area where SkillUp AI showed promise was in suggesting career paths. It is possible that SkillUp AI did better at recommending career paths compared with training or jobs because there were fewer career paths for the recommendation engine to choose from, leading to better fits with the users' inputs. Similarly, the reduction in the number of choices may have also helped SkillUp AI to perform better compared with SkillUp Explore, as the tools available in SkillUp Explore, such as the search filter, are less helpful when there are fewer options. This difference in content could help mitigate some of the limitations of the chatbot while leaning into some of its strengths, such as personalization. It seems feasible that subsequent versions of the tool — especially if it primarily targeted users interested in exploring career paths — could produce results similar or even superior to those of SkillUp Explore, though the value of such effort is questionable given the existence and familiarity of SkillUp Explore.

Looking Forward

AI-based chatbots have continued to grow in prominence and have shown promise in answering many basic questions and providing assistance in straightforward tasks. This experiment, however, suggests that a chatbot alone is probably not sufficient to help people navigate complex, high-stakes, nuanced journeys like career navigation. Given this reality, organizations like SkillUp need to continue to learn and adapt their AI strategies.

Moving forward, SkillUp will use the lessons from this experiment to continue developing AI tools for its website. This test showed that SkillUp AI as a stand-alone product is not currently a feasible replacement for SkillUp Explore. SkillUp now plans to provide an AI experience within SkillUp Explore — letting users ask the chatbot questions on the webpage while they are exploring a topic. For example, a user looking at a career page would be able to click a button to ask, “Does this career fit me?” and then engage with a chatbot to get personalized insight. SkillUp is also focused on building AI solutions that help users act, rather than just answering questions. SkillUp is identifying points where users often stop engaging with the website and where an AI tool could provide real-time support. The resulting experience is likely to look more like a summary of next steps (generated by AI) than a chatbot, bringing relevant content to the surface and adapting dynamically to what is known about the user.

Through the end of 2025, SkillUp plans to continue experimenting with AI tools, focusing on tools integrated into SkillUp Explore that address specific use cases (such as the examples described above), with evaluation assistance from MDRC. This rapid testing will allow SkillUp to learn what resonates best with users. In the long term, SkillUp will also observe whether a full, stand-alone chatbot might still add to its site, or, with the continued evolution of AI, LLMs, and consumer preferences, eventually fully replace SkillUp Explore.

Conclusion

The launch of ChatGPT in 2022 created a massive technological shift. In the context of billions of dollars in investment and claims of an oncoming AI revolution, expectations for LLMs were incredibly high.⁷ Rollouts for these tools have had mixed success, with some uses catching on, such as help in coding or transcription, while other applications have been less helpful or even harmful.⁸ These issues have shown the limitations of chatbots, and have contributed to a feeling that LLMs may not be as useful as once thought.

This pattern, termed the Gartner hype cycle, suggests that some innovations fail to meet the inflated expectations placed upon them during their initial prominence.⁹ However, after this period of “disillusionment,” many new technologies are slowly adopted and find productive uses, often proving indispensable once the technology has matured. Given the promise of LLMs, many companies have continued to develop AI tools and find useful ways to integrate AI into their offerings.

This test of SkillUp AI shows that AI-based technology targeting job seekers has not yet matured into the productive form suggested by the hype cycle. That fact does not mean that future development should not continue, however. While SkillUp AI did not outperform SkillUp’s existing platform in this A/B test, the test provided critical insights into how generative AI can and cannot support job seekers. This evaluation demonstrates the value of rapid testing and honest learning in emerging tech. It is hoped that other organizations will be able to benefit from this early attempt at applying this new technology to help job seekers. Future adaptation should help move these tools toward helping job seekers find opportunities that match their interests. The findings have directly informed SkillUp’s next steps, which should help future tools be more targeted, responsible, and effective for the users who need them most.

Notes and References

1. For an overview of evidence on career coaching, see Joseph Fuller, Kerry McKittrick, Sherry Seibel, James Wilson, Vasundhara Dash, and Alexandra Epstein, *Unlocking Economic Prosperity: Career Navigation in a Time of Rapid Change* (Harvard Kennedy School, 2023).
2. Sarah Shaw and Abigail Lupi, “Evaluating Career Development Outcomes for Coach, a GenAI Project by CareerVillage.org” (CareerVillage, 2024).
3. This approach is known as “retrieval-augmented generation.” See Amazon.com, “What Is RAG (Retrieval-Augmented Generation)?” (website: <https://aws.amazon.com/what-is/retrieval-augmented-generation>, n.d.).
4. Initially, SkillUp and MDRC planned to randomly assign 10,000 users in the A/B test, with 5,000 in each group. After examining the early impacts on engagement, it became clear to both teams that this version of SkillUp AI was not having positive impacts, and enrollment was stopped after 5,000 participants had been randomly assigned. Further analysis confirmed the initial assessment.
5. A Pew Research survey in early 2025 showed that only one-third of American adults had used an AI chatbot, while a YouGov poll found that the words most people were most likely to use to describe AI were “cautious,” “concerned,” and “skeptical.” See Colleen McClain, Brian Kennedy, Jeffrey Gottfried, Monica Anderson, and Giancarlo Pasquini, “How the U.S. Public and AI Experts View Artificial Intelligence,” (Pew Research Center, 2025) and Jamie Ballard, “Americans Are Increasingly Skeptical About AI’s Effects,” YouGov, 2025). Research into using AI chatbots for specific purposes and sectors has also found significant portions of the population are unwilling to use chatbots, mostly due to concerns about their accuracy. See Philip Fei Wu, Charlotte Summers, Arjun Panesar, Amit Kaura, and Li Zhang, “AI Hesitancy and Acceptability — Perceptions of AI Chatbots for Chronic Health Management and Long COVID Support: Survey Study,” *JMIR Human Factors* 11 (2024): e51086.
6. Most users spend less than 10 minutes on the platform, so even short delays of 30 to 45 seconds are meaningful.
7. Mike Isaac and Erin Griffith, “OpenAI Is Growing Fast and Burning Through Piles of Money,” *New York Times* (Sep. 27, 2024).
8. See Rhiannon Williams, “Why Google’s AI Overviews Gets Things Wrong” (website: <https://www.technologyreview.com/2024/05/31/1093019/why-are-googles-ai-overviews-results-so-bad/>, 2024); Jonathan Allen, “New York City Defends AI Chatbot that Advised Entrepreneurs to Break Laws” (website: <https://www.reuters.com/technology/new-york-city-defends-ai-chatbot-that-advised-entrepreneurs-break-laws-2024-04-04/>, 2024).
9. “Welcome to the AI Trough of Disillusionment” (website: <https://www.economist.com/business/2025/05/21/welcome-to-the-ai-trough-of-disillusionment>, 2025).

Acknowledgments

The authors would like to thank the staff members from SkillUp who collaborated with us on this project, collected and provided data, and shared their thoughtful contributions to the development of this brief. In particular, we thank Haleigh Boulanger, Desiree Jewell, Steve Lee, and Evan Marrow. We also thank our current and former MDRC colleagues Richard Hendra, who helped craft the vision for the brief and provided overall guidance and support; Cynthia Miller, who provided valuable comments on the brief; Mary Bambino, who managed the project budget and timeline; Alexandra Ubalijoro, who tested the chatbot and provided insights on its use; Joshua Malbin, who edited this brief; and Carolyn Thomas, who prepared it for publication. Finally, we thank the SkillUp users who participated in the study and allowed us to learn from their experiences.

The research reported here and the development of SkillUp AI were supported by the Gates Foundation, Cognizant, the Schultz Family Foundation, the Charles Koch Foundation, and others through grants to SkillUp Coalition.

The following organizations support dissemination of MDRC publications and our efforts to communicate with policymakers, practitioners, and others: Arnold Ventures, Ascendium Education Group, Yield Giving/MacKenzie Scott, and earnings from the MDRC Endowment. Contributors to the MDRC Endowment include Alcoa Foundation, The Ambrose Monell Foundation, Anheuser-Busch Foundation, Bristol-Myers Squibb Foundation, Charles Stewart Mott Foundation, Ford Foundation, The George Gund Foundation, The Grable Foundation, The Lizabeth and Frank Newman Charitable Foundation, The New York Times Company Foundation, Jan Nicholson, Paul H. O'Neill Charitable Foundation, John S. Reed, Sandler Foundation, and The Stupski Family Fund, as well as other individual contributors.

The findings and conclusions in this report do not necessarily represent the official positions or policies of the funders.

For information about MDRC and copies of our publications, see our website: www.mdrc.org.

Copyright © 2025 by MDRC®. All rights reserved.

New York
200 Vesey Street, 23rd Fl.
New York, NY 10281
Tel: 212 532 3200

Washington, DC
750 17th Street, NW
Suite 501
Washington, DC 20006

Oakland
475 14th Street, Suite 750
Oakland, CA 94612
Tel: 510 663 6372

