Unpaid referral program to drive GetCTC returns

Date implemented: July-August 2022
Date written: February 2023

Abstract: In July and August, GetCTC piloted a small referral program, in which 1,231 GetCTC clients with recently accepted returns received a message encouraging them to tell their friends and family about the tool. The results are inconclusive but suggest that the campaign probably generated a modest number of returns. Direct evidence (which is likely a significant undercount) suggests the generation of five returns, while regression results suggest the generation of zero to 115 returns, depending on specification. Estimates in the range of 20-50 seem plausible. While small in absolute terms, this would be a conversion rate far higher than most other text message campaigns, with one return generated per 20-60 messages sent; other relatively high-performing campaigns often generate one return per several hundred messages sent. Based on these results, we implemented a program of sending such referral messages on a weekly cadence to GetCTC clients who had successfully filed returns. The program sent 33,414 messages and—according to direct evidence, which is again likely to be a significant undercount—generated 188 submissions (178 messages/submission), of which 88 were accepted (380 messages/acceptance). While these results are not game-changing, they are highly effective on a per unit basis compared to other tax benefits outreach strategies from 2021 and 2022, and they were generated from a very modest referral program. Investing in and iterating on a referral program should be a high priority in outreach efforts for any future simplified filing or low-income filing tool.

Other experimental results and research from GetCTC 2022 are available here.

---

1. Background and research questions

Evidence in 2021 suggested that perhaps as many as one in four GetCTC users had heard about the tool from friends or family, although Code for America and our partners had done no proactive work to encourage and facilitate such referrals. (See Lessons from Simplified Filing 2021, p. 51.) In 2022, we chose to experiment with a GetCTC referral program, in an attempt to increase the number of referrals coming through this mechanism.

The project was initially designed as a paid referral program, akin to private sector analogues, where users would receive a payment for each person they successfully referred. Ultimately, the paid program was canceled due to legal and logistical concerns. Notably, the lack of financial incentives might not only decrease the likely effectiveness of the program, but also reduce the likelihood that we

would be able to detect any true effect. Referrers in a paid program could be required to use unique tracking URLs (at pain of not receiving their payments) whereas those in an unpaid program, even if they did alert friends and family, might not do so via any particular links, jeopardizing our ability to detect the program's effect.

Competing product priorities also limited the possibilities of the referral program. Higher-intensity interventions—like social media share buttons in the app upon submission of a return, or more hands-on encouragement of referrals—were not possible given limited capacity.

In July and August, we ran a very small-scale pilot referral program without financial incentives. The pilot was designed both to limit engineering effort and to increase the probability that an effect could be detected even if referrers did not use unique tracking URLs. Given the lack of financial incentives, the speculative experimental design, and the limitation on engineering resources, a null effect should not be interpreted as indicating that referral programs cannot be successful—simply that we did not detect an effect from this limited program.

2. Study design and implementation

On two Thursdays in July and August, GetCTC clients whose returns had been accepted in the last five days received a message through their preferred form of communication (text or email) in their preferred language (English or Spanish), reading:

> "Thank you for filing your tax return with GetCTC! Every family should get the money they deserve. Tell your family and friends to claim their benefits too at getctc.org/refer"

Returns could be directly tracked if referrers were to forward the full getctc.org/refer link to their friends and family, and the recipients clicked it—though previous projects suggested referrers were unlikely to send the complete link and would instead simply refer people they knew to the generic getctc.org home page.[1] To increase the likelihood of detecting an effect without detecting it from unique URLs, we introduced geographic randomization into the program. On July 28, clients in 25 randomly selected states received messages, shown as Group 1 below. On August 4, clients in the other 25 states received messages. If referrals are somewhat more likely to remain in-state than out-of-state—a plausible but not air-tight assumption—then the difference in submission rates between Group 1 and Group 2 states in the week between July 28 and August 4 should be a portion of the effect attributable to the program. Moreover, if most of the effect of the program occurs within a week—which is consistent with the timing effects of other interventions—then the difference in rates between the groups of states the week after August 4 would also be attributable to the program.

[1] Note that when a client arrives at the getctc.org/refer link, the application stores the "refer" parameter and redirects to the URL getctc.org. This means that if a messaging recipient clicked the link, copied the URL from their browser bar, and pasted it to a friend, the link they sent the friend would not have the "refer" parameter attached.

	Group 1 states (7/28)	Group 2 states (8/4)
States	AL, AK, AZ, CO, CT, DE, GA, HI, ID, IN, KS, LA, MI, NH, NJ, NM, NY, ND, OH, SC, SD, TN, TX, UT, VA, WV	AR, CA, DC, FL, IL, IA, KY, ME, MD, MA, MN, MS, MO, MT, NE, NV, NC, OK, OR, PA, RI, VT, WA, WI, WY
Total messages sent	503	728

3. Results

3.1 By unique URL

As expected, very few clients in absolute terms actually filed returns through the /refer unique link; only four submitted returns, all of which were accepted. (A fifth began a return but was flagged for ID verification and did not finish submitting.)

Unique home page views	179
Started flow	25
Cleared screeners	14
Started return	7
Submitted returns	4
Accepted returns	4 (all RRC only)

While these numbers are small in absolute terms, the conversion rate is reasonably high in context. Even if there were no referrals that went through other unique URLs, the campaign would have a conversion rate of 308 outgoing messages per accepted return—far better than, for example, messages from state agency partners, which sometimes have rates of 1,000 outgoing messages per accepted return.

3.2 By geography

Recall that results will only show up geographically if (1) filers are more likely to refer other clients *in their states* and (2) they are likely to do so within a week of message sending. We can use data from the unique URL usage to partially test these theories—and we find that they do not entirely hold out. The five started returns using the /refer URL were started on the following days:

State	Date message sent in that state	Date return started
TX	7/28	7/28
NM	7/28	7/30

CA	8/4	8/4
MI	7/28	8/7
MI	7/28	8/12

While the first three returns follow the predicted patterns, the last two do not; these intakes were started 10 and 15 days after the referral messages were sent in Michigan. As such, it is far from clear that we should expect to see any clear signal in the geographic results.

With this caveat in mind, the table below shows the impact of the referral program in the week after messages were sent. Regressions are OLS at the level of the state, with returns in the baseline week as a control. Columns (1), (3), and (5) show results on Week 1 returns, viewing the Group 1 states as the treatment group; columns (2), (4), and (6) show results on Week 2 returns, viewing the Group 2 states as the treatment group. The dependent variable is submitted returns in the given week.

	All returns		Dropping fraud		Dropping fraud; null source	
	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Week 1	Week 2	Week 1	Week 2	Week 1	Week 2
Treated group	-9.052	-4.562	2.426	7.751	1.866*	2.676***
	(12.13)	(27.31)	(4.521)	(15.22)	(1.037)	(0.816)
Count baseline wk	0.943***	1.134***	0.947***	0.909***	0.844***	0.792***
	(0.0530)	(0.119)	(0.0213)	(0.0716)	(0.0570)	(0.0449)
Constant	25.05	39.49	-1.298	2.334	-3.350*	-4.270***
	(19.09)	(42.99)	(7.125)	(23.99)	(1.688)	(1.328)
Observations	51	51	51	51	51	51
R-squared	0.869	0.654	0.977	0.774	0.822	0.870

Standard errors in parentheses
* = p<.1, ** = p<.05, *** = p<.01

In columns (1) and (2), we look at results on all returns. In columns (3) and (4), to reduce noise, we drop likely-fraudulent returns as determined by our fraud filters and ex-post review processes. In columns (5) and (6), to further reduce noise, we restrict the analysis to non-fraudulent returns *with a null source code*. These represent a little over 10% of returns. In general, a null source indicates that the return was not the result of a different outreach campaign, and so we should expect most friend-to-friend referrals to have a null source. (This assumption is not guaranteed to hold. It is possible, for example, that a referral from a friend could lead a client to google GetCTC and click on a digital ad, which would have a source code; or the referral could lead a client to look into GetCTC online, find ChildTaxCredit.gov, and arrive from that site, which would have a source code.)

While results in columns (1) through (4) show no treatment effect, the results in columns (5) and (6) are significant, indicating that the referral program generated about two extra returns per state in the week after messages were sent. Taken at face value, the point estimates would suggest that the program generated 115 returns—an estimate that seems unlikely on its face given the null results in columns (1) through (4).

Appendix A shows results at the aggregate level across all states—by count and graphically—and does not show any such effect.

3.3 By self-report of referral source

A final way to assess the referral program's impact is to look at self-reports of friend/family referrals. After a client successfully completes a GetCTC return, they receive a message asking them to complete a follow-up survey about their experience, which includes a question asking how they heard about the tool. "From friends/family" is one of the possible answers. Of course, clients may misreport how they heard about the tool, and only a small and potentially unrepresentative sample of clients take the survey. The figure below shows the fraction of survey respondents who provided this answer over time, with the two dates of referral messages being sent shown in red. One should expect any effect to show up at a bit of a lag from the date of the messages being sent. The results do not suggest any particular effect of the referral program.

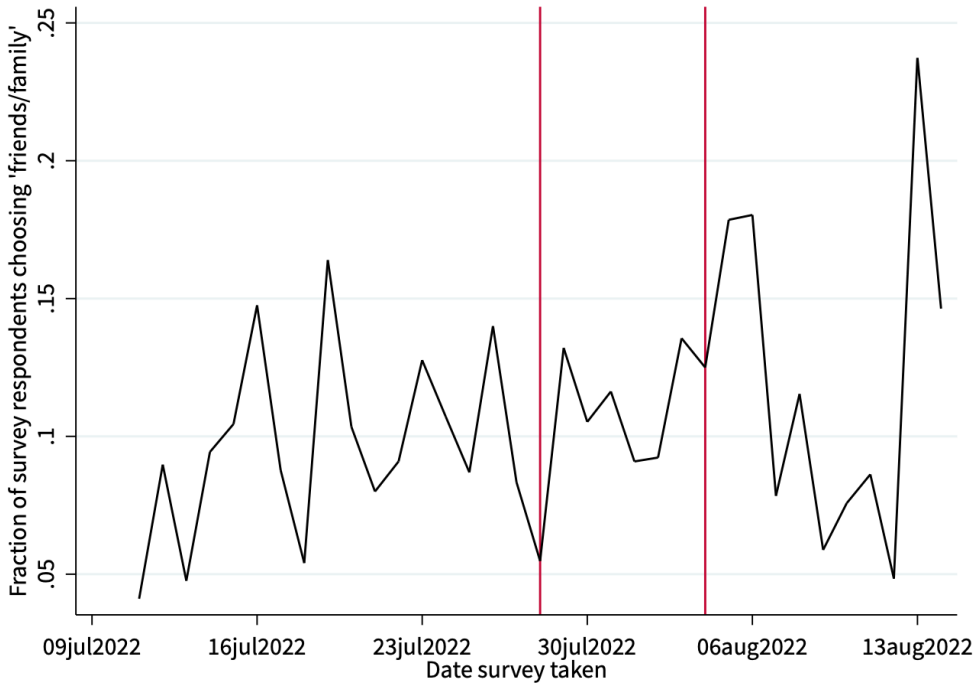

4. Results of follow-up referral program

Based on the results of this pilot, GetCTC implemented such a program throughout the last few months of GetCTC operation in 2022. Beginning on September 26, and weekly on Mondays thereafter, we sent messages—via text message or email, and in English or Spanish, according to the client's preference—to clients who had submitted returns in the last week that were accepted by the IRS. As in the pilot, the messages encouraged recipients to refer their friends to getctc.org/refer. Over seven weeks, the program sent 33,414 messages and saw:

Unique home page views	8.93% (2,985)
Started flow	2.59% (866)
Cleared screeners	1.34% (447)
Started return	1.06% (354)
Submitted returns	0.56% (188)
Accepted returns	0.26% (88)

Again, these numbers—which are surely an undercount, given the nature of the unique URL tracking—are somewhat low in absolute terms but relatively quite high. It took 178 outgoing messages to generate one submission, and 380 outgoing messages to generate one accepted return. These numbers are roughly in line with the conversion rate from the pilot.

5. Discussion

Trusted messengers with a personal connection are invaluable in countless settings, and motivating low-income households to file a tax return is likely no different. This paper reported results from a limited study of a very modest referral program that nevertheless generated cost-effective results. In the pilot phase, sending a single follow-up text message to clients who had already used our service generated at least one GetCTC submission for every 246 messages sent, and perhaps one submission for as few as every 25 messages sent. The follow-up implementation of this program generated at least one GetCTC submission for every 178 messages sent, which is, methodologically, certainly an undercount.

These results are especially striking in the broader context, in which effective outreach interventions to promote tax benefits to very low-income households have been hard to find. Many outreach interventions took thousands of more-expensive contacts to generate a submission; even some of the best methods took hundreds per submission and a thousand per accepted return. (See Lessons from Simplified Filing 2021, Part 3.3.)

Recall, too, that this was a very low-touch referral program, designed primarily based on the legal and logistical limits of the GetCTC program, and with a minimum of user research or iteration. While it is

possible that we stumbled into an optimal program, it is unlikely; it is more likely that further iteration and additional resources could produce a still more effective program. Such an effort ought to be a priority outreach strategy for any simplified filing or low-income filing tools in the future.

Appendix A. Aggregate returns by treatment group

	7/21-7/27 (baseline)	7/28-8/3 (Group 1 treatment period)	8/4-8/10 (Group 2 treatment period)
Panel A: All returns			
Group 1 submitted (all)	2,071	2,381	3,214
Group 2 submitted (all)	2,814	2,972	3,928
Total (all)	4,885	5,353	7,142
Fraction Group 1 (all)	.4240	.4448	.4500
Panel B: Not including likely fraud returns			
Group 1 submitted	1,900	1,854	1,897
Group 2 submitted	2,559	2,659	2,649
Total	4,459	4,513	4,546
Fraction Group 1	.4261	.4108	.4173
Panel C: Not including likely fraud returns; null source only			
Group 1 submitted	241	196	170
Group 2 submitted	266	275	268
Total	507	471	438
Fraction Group 1	.4753	.4161	.3881

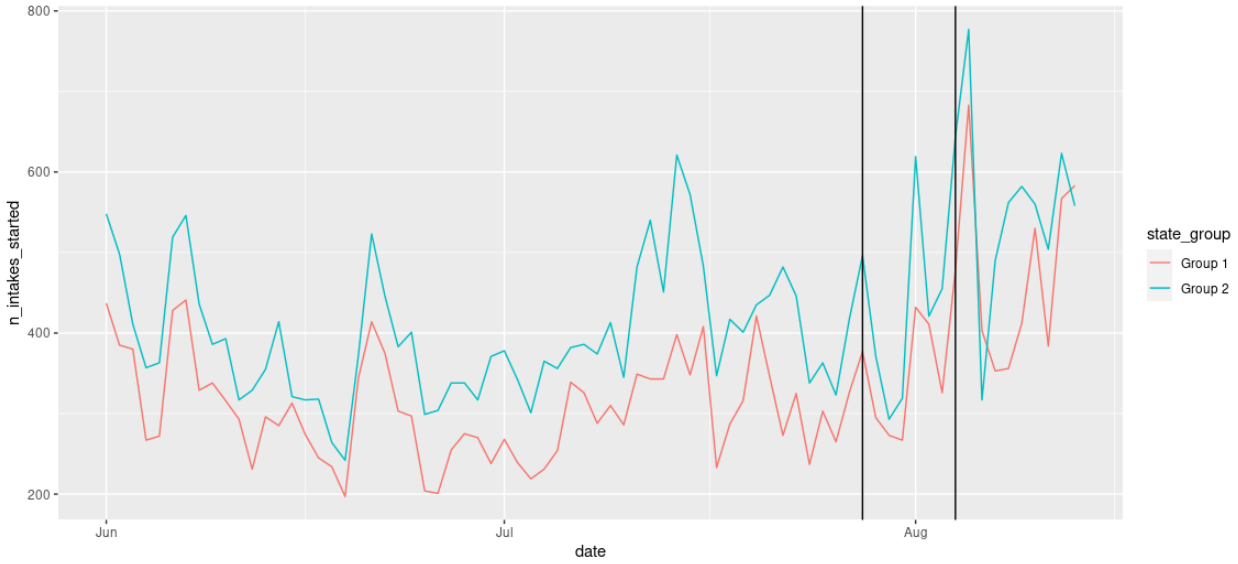
800
600
400
200
n_intakes_started
Jun
Jul
Aug
date
state_group
Group 1
Group 2