

Digital Marketing Interim Evaluation: Intervention Cycle 2

Grantee Agency: Virginia Division of Child Support Enforcement
Grant Number: 90FD0218-01-01
Date of Submission: May 1, 2020
Title of Intervention: Find/Engage/Educate
Key Staff Contact Information: Mariellen Keely
Policy Program Consultant - Initiatives
Virginia Department of Social Services
801 East Main Street, Richmond VA 23219
mariellen.keely@dss.virginia.gov
(804) 726-7923

Contents

- Executive Summary..... 1
- 1. Background 4
 - I. Grant Purpose..... 4
 - II. Problem 4
- 2. Intervention 5
 - I. Goals of the Overall Project..... 5
 - II. Goals of Intervention Cycle 2 (Engage)..... 6
 - III. Cycle 2 Research Questions..... 7
 - IV. Development of Cycle 2 Intervention 8
 - A. Revising Initial List of Activities for Cycle 2 8
 - B. Selecting Social Media Channels for Organic and Paid Campaigns 9
 - C. Developing and Maintaining a Dynamic Budgeting Model for the Paid Campaign..... 10
 - D. Determining the Messaging Focus for the Paid and Organic Campaigns 11
 - E. Determine Approach to Geographic Targeting for Ads in the Paid Campaign 12
 - F. Determining the Types of Images to Be Used in the Paid and Organic Campaigns 14
 - G. Determining Approach to Video in the Paid Campaign 15
 - H. Creating Content Calendars for the Organic and Paid Campaigns 15
 - I. Registering Social Media Accounts on Chosen Platforms..... 17
 - J. Ensuring DCSE’s New Online Child Support Application Is Operational 17
 - V. Cycle 2 Description 19
 - A. Creating a Website to Track Application Conversions..... 19
 - B. Developing, Executing, and Monitoring the Organic Social Media Campaign 21
 - C. Developing, Executing, and Monitoring the Paid Social Media Campaign 24
 - D. Developing and Executing a Social Media Hashtag Campaign 35
 - E. Revising the Budget and Content Calendars in Response to Changes in Organic and Paid Campaigns..... 36
 - F. Responding to Comments on Posts in the Organic and Paid Campaigns 36
 - VI. Target Populations 39
 - VII. Timelines..... 40
 - A. Overall Project Timeline..... 40
 - B. Detailed Timeline for Cycle 3 40

VI. Cycle 2 Outcome Measures.....	41
3. Results.....	41
I. Sample Size.....	41
II. Results.....	42
III. Analysis.....	44
A. The Organic Social Media Campaign.....	44
B. Comments and the Hashtag Campaign.....	45
C. The Paid Social Media Campaign.....	45
D. Online Child Support Application.....	55
4. Lessons Learned and Next Steps.....	56
I. Lessons Learned.....	56
II. Next Steps.....	57
Schedule of Appendixes.....	59
Appendix A: Apply Webpage.....	60
Appendix B: Lead Form Ad.....	61
Appendix C: Ad Performance in Cycle 1 versus Cycle 2 (Partial), Ranked by Spends in Top-Ten Designated Market Areas (DMAs).....	62
Appendix D. Contact Form Responses.....	64

Executive Summary

This report offers an interim evaluation of the second cycle of interventions undertaken by Virginia’s Division of Child Support Enforcement (DCSE) as part of the Digital Marketing grant program, sponsored by the federal Office of Child Support Enforcement. These interventions took place during the 90-day period beginning at 12:01 a.m. on November 1, 2019 and ending at midnight on January 29, 2020.

After explaining the purpose of the grant, the problem we intended to address through this intervention, and our research questions, we describe the six major components of this cycle (ordered according to roughly when the associated process began):

1. Creating a webpage to track anonymized records of clicks to a newly created child support application website. These clicks would be treated as conversions, in the advertising sense.
2. Developing, executing, and monitoring an organic social media campaign¹ that targeted parents and other child support related audiences directly. We wanted to see whether the immediacy and relatively low cost of an organic campaign could make it easier to disseminate up-to-date program information as well as give program participants a chance to ask questions or seek help with a new or existing child support case. We also anticipated receiving negative comments and wanted to use this time to understand how best to respond to those.
3. Developing, executing, and monitoring a paid social media campaign that would complement the organic media campaign and perhaps build on the electronic advertising undertaken in Cycle 1. We drew text and images tested in that earlier cycle but added short videos that used cartoon-style animation to convey this cycle messaging. At the same time, we planned to use a more fully articulated funnel approach to our advertising, described in section 2.IV.D. below, that was intended to move users from an initial state of awareness through to a significant final action: applying to open a child support case.

¹ In this report we follow the common contemporary marketing practice of using the word “organic” to distinguish marketing approaches that platforms are not paid to promote. In practice, as we note below, organic content that seemed to receive unusually high levels of engagement (as demonstrated by, for example, the number of times users liked or shared that content) was sometimes promoted through advertising funds so it could reach a wider audience. This specific part of our organic campaign thus became part of our paid campaign.

4. Developing, executing, and monitoring a social media hashtag campaign² to unite the organic and paid campaigns and hopefully encourage other users to share positive content related to child support in Virginia.
5. Revising the budget and content calendars in response to changes in organic and paid campaigns.
6. Monitoring and responding to social media comments.

As we describe in more detail below, we implemented all six components of the intervention, though not all aspects of each component were implemented as planned. During Cycle 2:

1. We created our planned webpage to promote DCSE’s newly developed online child support application, as discussed in section 2.V.A., and recorded 182 conversions from the webpage to the application, as noted in section 3.III.D.
2. We created organic social media accounts and led related organic campaigns on Facebook, Instagram, and Twitter, albeit with some difficulty, as detailed in section 2.IV.I.
 - We made 24 organic posts on each of these platforms, as noted in section 2.IV.B.II.
 - By the end of Cycle 2, DCSE’s Facebook account had approximately 200 followers, its Twitter account had approximately 90, and its Instagram account had more than 500, as noted in section 3.III.A.
3. Our paid campaign released at least 149 advertisements across the three platforms, as noted in section 2.V.C.IV, at a cost of just under \$22,000
 - We fell below our goals in Cycle 2 for both Facebook and Instagram, if measured using the same type of click-through rate used in cycle 1. In Cycle 2, our click-through rates for these platforms—or rather, strict click-through rates—were 26 percent below the goal for Facebook and 25 percent below for Instagram. These results are discussed below in sections 3.II. and 3.III.C.

² Used across many social media platforms, hashtags are words or phrases that allow users to group otherwise unrelated content across that platform. It typically begins with a pound or hash sign and does not include spaces. A guide to this kind of campaign can be found here: <https://www.tintup.com/blog/how-to-use-hashtags-for-campaigns/>.

- On Twitter we achieved a click-through rate of 5.44 percent, putting us 52 percent above our goal of a click-through rate of 2.62 percent for that platform. These results are also discussed below in sections 3.II. and 3.III.C.
 - Combined, our organic and paid social media campaigns reached approximately 536,000 unique users across Facebook, Instagram, and Twitter, as noted in section 3.I.
 - We saw some evidence that social media advertising and DCSE’s new online child support application—perhaps in combination or perhaps primarily because of the effect of one or the other—was associated with a less pronounced decline in new applications for child support services in Virginia. In other words, the overall number of new applications during Cycle 2 fell somewhat from the same period during the previous federal fiscal year but that fall was not as precipitous as it had been in previous years. While not the same as meeting our overall goal of increasing the absolute number of applications from the same time during the previous year, this tentative evidence of does seem to offer a practical possibility for other programs seeking to mitigate declining child support caseloads.
4. We launched a hashtag campaign (#supportVAkids) and placed that hashtag in virtually all of our organic and paid posts. Yet, as discussed in sections 2.V.D and 3.III.B., we did not see wider support for the hashtag outside of our own posts or those of individuals or organizations directly related to this project.
 5. We made frequent revisions to our content calendars and budgeting model in an attempt to improve our results or address problems. Tracking those changes and ensuring that they were executed properly took a fair amount of time for project staff and represents one of the more significant challenges for running a complex set of social media campaigns.
 6. We received and responded to a small number of negative comments on our organic and paid posts, as analyzed in sections 2.V.F. and 3.III.B.

We conclude with a discussion of lessons learned and next steps. A set of four appendixes provides readers with additional data and examples of materials that the project team developed for this cycle.

1. Background

I. Grant Purpose

The Digital Marketing grant program, sponsored by the federal Office of Child Support Enforcement (OCSE) within the U.S. Department of Health and Human Services' Administration for Children and Families, is a 24-month demonstration project with the goal of researching how digital marketing may help the child support program more effectively reach and serve families. In September 2018, OCSE awarded funds to 14 child support agencies to test digital marketing approaches and partnerships to reach parents that could benefit from child support services, and create or improve two-way digital communication and engagement with parents.

II. Problem

The IV-D program in Virginia faces a challenge common to other programs across the country: how to bridge the gap between the decline in the number of cases in its caseload and the ongoing need for its services demonstrated through analyses of population data. This intervention addresses the problem of how to bring these two sets of data points into better alignment.

Between federal fiscal years 2013 and 2018 the number of total child support cases in Virginia with either current or past child support due declined by 10.8 percent, or more than 35,000 cases. Nearly 20,000 of those cases (or 55 percent of the total decline) came from Former Assistance. Cases classified as Never Assistance, on the other hand, fell by less than 5 percent (approximately 6,000 cases) over the same time, suggesting that there remains a relatively robust and ongoing need for IV-D services among this portion of the caseload.

An analysis of Census data and caseload data from the Virginia Division of Child Support Enforcement (DCSE) gives a sense of how significant that need might be. According to Census data for 2015, approximately one in four children nationwide had at least one parent out of the home and appeared to qualify for child support services.³ Caseload data for a recent comparable period, Virginia's 2016 fiscal year, showed that approximately one in five Virginia children were receiving IV-D services from DCSE. Presuming that the nationwide figure is roughly representative of Virginia, the difference between the overall need and the caseload data is approximately 93,000 children. Adjusted for Virginia's average of 1.2 children per child support case, this yields a potential increase of approximately 77,000 new child support cases, or 26 percent, over 2018 levels.

³ T. Grall (2018). *Custodial Mothers and Fathers and Their Child Support: 2015*. U.S. Census Bureau, Current Population Reports. <https://www.census.gov/content/dam/Census/library/publications/2020/demo/p60-262.pdf>.

More current U.S. Census data suggest that the gap between families eligible for services and those participating in the IV-D program may be widening. According to American Community Survey (ACS) estimates, there were between 540,416 and 561,162 children under the age of 18 in single-parent households in Virginia in 2017.⁴ DCSE caseload data, on the other hand, identified 301,284 children under 21 in the DCSE caseload in federal fiscal year 2017. That leaves a minimum of roughly 240,000 children in Virginia who could presumably be eligible for IV-D services.⁵

Based on market research conducted by other IV-D programs as well as broader scholarship on the public perceptions of child support,⁶ we hypothesize that the reason for this gap in service is that many members of the public either remain unaware of the federal child support program or have significant misconceptions about it. We believe that digital engagement methods can increase awareness about both the program's existence and how it actually delivers services to families in Virginia and elsewhere.

DCSE leadership—particularly Craig M. Burslem, Deputy Commissioner of State Programs for the Virginia Department of Social Services (VDSS)⁷—was largely responsible for defining and prioritizing the project team's focus on increasing program participation, particularly among Never Assistance families, though their collective awareness of declining caseloads was also informed by other discussions and research in the child support community as a whole.

2. Intervention

I. Goals of the Overall Project

The primary purpose of the overall project is to gather information about possible strategies to close the apparent gap, described above in section 1.II.A, between the number of families eligible for services and those currently connected to Virginia's IV-D program. To accomplish that, Virginia seeks to increase requests for child support services, with a focus on applications for the core child support services (locating parents, establishing paternity and child support orders, collecting and distributing child support, and enforcing and modifying child support orders) from custodial parents who are eligible for child support services but not currently

⁴ U.S. Census Bureau, 2013-2017 American Community Survey 5-Year Estimates.

⁵ Even this large number may understate the need. Among other factors, the mismatch in the age ranges between the two measures (under 18 for ACS data, under 21 for DCSE data) would necessarily mean there are actually even fewer children under 18 on the DCSE caseload.

⁶ Perhaps the best-known study of public perceptions of child support was conducted by the public relations agency Ogilvy on behalf of California's IV-D program in 2018.

⁷ At the time of the original grant application, Deputy Commissioner Burslem was the Director of DCSE.

connected to the IV-D program and who have never received public assistance benefits (Never Assistance).⁸

To achieve this goal, we mapped a three-part intervention that reflects a somewhat simplified version of the Transtheoretical Model of Change, originally described by James O. Prochaska and Carlo DiClemente.⁹ While their model described five overall stages (Precontemplation, Contemplation, Preparation, Action, Maintenance, and Relapse), we focused primarily on the three middle sections (Contemplation, Preparation, and Action).

Looking at those stages from the outside—that is, from the perspective of a program such as ours, rather than from the individual decision maker described in the model—we recast those sections into three intervention cycles, which we named Find (Cycle 1), Engage (Cycle 2), and Educate (Cycle 3).

An interim evaluation for DCSE’s intervention during Cycle 1 (Find) was approved on March 30, 2020. The focus of this report is Cycle 2 (Engage).

II. Goals of Intervention Cycle 2 (Engage)

The overall goal of second cycle of interventions is the same as that of the overall project: encouraging more families to apply for services IV-D program, as measured by the number of applications. We approached that in cycle 2 by attempting to broaden public awareness of DCSE and enrich the program’s engagement with new and existing clients. In particular, we focused on increasing awareness of Virginia’s newly created online child support application and encouraging users to apply for services through it.

In planning and executing this cycle, we relied on the Learn, Innovate, Improve (LI²) paradigm to shape our plans and structure our approach to modifying our current and future output.¹⁰ We saw the LI² paradigm as particularly important for helping us understand the effectiveness of

⁸ We use “public benefits” here to refer to the small number of assistance programs, primarily TANF, that distinguish child support case types for purposes of reporting on Form OCSE-157.

https://www.acf.hhs.gov/sites/default/files/programs/css/at_14_09b.pdf.

⁹ J. O. Prochaska and C. DiClemente (1983). Stages and Processes of Self-Change of Smoking: Toward an Integrative Model of Change. *Journal of Consulting and Clinical Psychology* 51(3): 390–5.

https://www.researchgate.net/publication/16334721_Stages_and_Processes_of_Self-Change_of_Smoking-Toward_An_Integrative_Model_of_Change. For a fuller discussion of how the Transtheoretical Model of Change relates to our approach, see our original grant application (“Virginia Child Support Digital Marketing Demonstration”).

¹⁰ More information about the LI² paradigm is available in the original OCSE funding announcement (“Using Digital Marketing to Increase Participation in the Child Support Program,” HHS-2018-ACF-OCSE-FD-1368, https://ami.grantsolutions.gov/files/HHS-2018-ACF-OCSE-FD-1368_0.pdf) and through 2017 Office of Planning, Research & Evaluation (OPRE) report “Learn, Innovate, Improve (LI²): Enhancing Programs and Improving Lives” (https://www.acf.hhs.gov/sites/default/files/opre/li2_brief_final_b508.pdf). OPRE Report #2017-108.

written and visual content and how those might be combined most effectively and adapted for reuse in Cycle 3 and beyond.

III. Cycle 2 Research Questions

To evaluate how successfully our interventions supported the goals of Cycle 2 (Engage), we sought to answer the following research questions:

1. What kinds of social media marketing strategies might drive positive public engagement with child support services?
2. Can promoting an online application portal through social media marketing increase the number of new applications for child support services?

Our view of digital engagement is informed by an analysis from the Internet Advertising Bureau (IAB). In a report titled “Digital Ad Engagement: An Industry Overview and Reconceptualization,” the IAB describes three distinct categories of engagement: cognitive, emotional, and physical.¹¹ The authors see cognitive engagement as encompassing mental states related to taking in and acting on information, particularly “awareness, interest, and intention,” that could be measured by such means as heart monitoring, eye tracking, and pupillometry (the measurement of pupil size and reactivity). Emotional engagement considers a user’s felt response. For some marketing techniques, this might be gauged through surveys or focus groups. For others, such as the type of social media activities conducted in Cycle 2, it could be assessed by looking at the number of likes a post received or the content of comments on those posts. Physical engagement is perhaps the most commonly measured type of engagement, since it can be recorded through clicks, shares, and conversions.

The IAB authors distinguish between two points that are often conflated: the intangible qualities of individual user experiences, on the one hand, and the methods used to capture evidence of those experiences, on the other. Maintaining this disinclination is especially valuable in the context of a formal evaluation, where metrics can sometimes seem to represent the whole of engagement rather than merely the visible evidence of it. We have tried to follow their approach in this report.

¹¹ Internet Advertising Bureau (2013). “Digital Ad Engagement: An Industry Overview and Reconceptualization.” <https://www.iab.com/wp-content/uploads/2015/05/IABAdEngagementWhitepaperDec2012FinalFinal.pdf>. As the report notes from the perspective of 2013, “Nearly two decades into the growth of online advertising, ‘engagement’ is still one of the most used, yet least understood terms [in advertising].” Little appears to have changed in the intervening period, and examples of online marketers and researchers trying to grapple with the meaning of “engagement,” and especially “digital engagement,” remain plentiful.

We should acknowledge that the bulk of this report addresses the first of our two research questions, since the activities of the intervention directly relate to that. We do offer data and some tentative analyses related to our second, largely exploratory research question (primarily in section III.3.D) and hope to be able to offer a fuller discussion of this in our final evaluation.

IV. Development of Cycle 2 Intervention

A. Revising Initial List of Activities for Cycle 2

Plans for Cycle 2 were initially developed by DCSE leadership as part of the original grant application. In that iteration this cycle focused on increasing engagement with parents by first engaging with judges in Virginia's Juvenile and Domestic Relations district courts, which control the judicial side of child support in the state. The plan had three components:

1. Conduct a focus group of Juvenile and Domestic Relations judges
2. Use the results of that focus group to alter any existing ad placements or to revise the DCSE website
3. Develop a print handout that includes a so-called short code that they would use to send a text a message to DCSE, which in turn would send them a link to DCSE's recently redesigned website

In February 2019 these plans were revised by the primary members of the digital marketing project team, which included DCSE project staff, the project's primary internal partner (the VDSS Public Affairs unit), and its outside vendor (Grays Peak Strategies). The consensus at the time was that introducing a print component into the outreach mechanism seemed at odds with the grant's overall focus on digital media. It also seemed harder to apply insights gleaned from an LI² based analysis of Cycle 1 to this initial vision of Cycle 2.

We also decided we should use Cycle 2 to promote DCSE's distinct but related development of an online child support application, since it aligned directly with our goals of increasing applications and creating an electronic (rather than paper-based) end point for user engagement.

As a result, the initial project team revised the plans for Cycle 2 to include six components:

1. Creating a website to track anonymized records of clicks to the newly created child support application website. These clicks would be treated as conversions, in the advertising sense

2. Developing, executing, and monitoring an organic social media campaign that targeted parents and other child support–related audiences directly (rather than through judges, for example)
3. Developing, executing, and monitoring a paid social media campaign
4. Developing, executing, and monitoring a social media hashtag campaign¹² to unite the organic and paid campaigns and hopefully encourage other users to share positive content related to child support in Virginia
5. Revising the budget and content calendars in response to changes in organic and paid campaigns
6. Monitoring and responding to social media comments

B. Selecting Social Media Channels for Organic and Paid Campaigns

Once our plans were set, we turned to the fundamental question of which social media platforms we would use for our organic and paid campaigns. From the beginning we were primarily drawn to Facebook, Instagram, Twitter, and Snapchat, largely because of their large user base¹³ and because their formats encouraged direct, two-way engagement, as opposed to sites such as YouTube or Spotify, which seem to emphasize professional or semi-professional content creation over less formal ways of building community.

During July and August 2019, while still in Cycle 1, we met with four professionals with experience administering social media platforms dedicated either to child support or paternity establishment. While their responses diverged on many points, all four agreed that Facebook and, to a lesser degree, Twitter could be effective ways to build awareness of the program and its services. Two respondents also recommended Instagram, with others unsure that the program’s focus on images over informational content was well-suited to child support.¹⁴ None

¹² Used across many social media platforms, hashtags are words or phrases that allow users to group otherwise unrelated content across that platform. It typically begins with a pound or hash sign and does not include spaces. A guide to this kind of campaign can be found here: <https://www.tintup.com/blog/how-to-use-hashtags-for-campaigns/>.

¹³ According to the data site Statista, Facebook had an average of 170 million monthly users on its mobile app, as of September 2019, while Instagram had 121 million, Twitter 81 million, and Snapchat 46 million. Also included on that list are Facebook Messenger (106 million), Pinterest (67 million), and Reddit (48 million), none of seemed well-suited to our campaigns. Statista, “Most Popular Mobile Social Networking Apps in the United States as of September 2019, by Monthly Users.” <https://www.statista.com/statistics/248074/most-popular-us-social-networking-apps-ranked-by-audience/>

¹⁴ Our own experience on Instagram could be said to reflect this same uncertainty, with our organic campaign finding some success there, thanks perhaps to its integration of information into the images themselves, while our paid campaign never quite found its footing. Data associated with these analyses appears below in section 3.III.A.

of the people we spoke to felt that Snapchat was likely to be a productive channel for organic or paid social media, outside of the promotion of events.

By the middle of September 2019, the project team, DCSE leadership, and VDSS Public Affairs staff had agreed that we would use three platforms in Cycle 2 for both organic and paid content delivery: Facebook, Instagram, and Twitter. The organic content would come from DCSE-specific social media accounts that we would launch on each of these three channels as part of the project.¹⁵ Indeed, one reason for selecting these three channels was to see whether paid advertising might be helpful in building DCSE's organic social media presence.

C. Developing and Maintaining a Dynamic Budgeting Model for the Paid Campaign

While working to finalize which social media platforms we would use in Cycle 2, we also developed a more complex budgeting tool to forecast and track expenditures across multiple platforms and ad types.

Our desire to develop this tool was in part a response to the somewhat rigid budgeting method we employed in Cycle 1. As we noted in our interim evaluation of that cycle:

We changed our Search budget several times over the course of the campaign, as we came to believe that the campaign was underfunded relative to demand. By contrast, the relatively poor return on investment offered by the Display campaign—at least as we executed it, without a full focus on remarketing and conversion targeting—made us wish that we had put less of our budget into that side of the intervention and potentially reserved it for later parts of the project.

We also felt obligated to take a more granular approach to budgeting in Cycle 2 to accommodate the greater variety of platforms and formats we anticipated using. In Cycle 1 all of our ads had been placed through a single company (Google) and took one of two forms: Search ads or Display ads.¹⁶ In Cycle 2 we anticipated using at least two companies (Facebook, which also owns Instagram, and Twitter) and imagined that we might use multiple formats (text and image ads, video ads, promoted posts, lead form ads, etc.). Each of these formats might need to have its own budget line.

Ultimately, we developed a spreadsheet that allowed us to make weekly budget adjustments in dollars to individual platforms on a weekly basis and, within each platform, to adjust the percentage of the overall weekly budget being allocated to a specific type of ad. Changes in the

¹⁵ Until shortly before the launch of Cycle 1, all of DCSE's social media messages were delivered by the agency-wide VDSS accounts on Facebook and Twitter; VDSS does not have an active presence on Instagram.

¹⁶ We capitalize "Search" and "Display" when referring to Google's proprietary distribution networks for those ad types. Lowercase "search" and "display" refer to those generic ad types, regardless of the platform used to distribute them.

weekly platform budget were automatically populated to the percentage budget in the ad-type area and both used conditional formatting and checksums to ensure that we stayed on budget.

D. Determining the Messaging Focus for the Paid and Organic Campaigns

Shortly after the end of Cycle 1 we began analyzing the outcomes of the three content areas that were represented on the main project webpage¹⁷ and targeted in our ads (applying for child support, paternity testing, and family engagement services, particularly co-parenting and employment-related services) to see whether we should continue to focus across all three or choose only one or two.

While individual ads for each of the three areas exceeded project goals in terms of click-through rates and related measures, users requested information about the three areas at significantly different rates. Approximately 55 percent asked to receive an application for child support services, 23.5 percent posed a question about an existing child support case, 7.7 percent asked about paternity services, and 5.3 percent about family engagement services; another 9.5 percent had questions or requests that fell outside any of these categories.

Looking at this data in preparation for Cycle 2, we came to see the relatively small percentage of requests related to paternity and family engagement services as evidence of comparatively low demand for those services—or at least for receiving them through DCSE. Other explanations are of course possible, including ones focusing on the language used to advertise or explain those services or related to the limitations of the main project website.¹⁸

Based on these data and the high likelihood that DCSE would be able to successfully create an online child support application ahead of the launch of Cycle 2, we decided that it made more sense to narrow our focus to only one area: encouraging users to apply for child support.

We also decided to structure the content in this more narrowly focused campaign in keeping with the widely used concept of the advertising funnel, which brings potential clients from Awareness to Intent and then Action. The similarities between this approach and the theory of change being used in our project as a whole is suggested in Figure 1.

¹⁷ In this report we use “main project webpage” to refer to the URL given in advertising as supportVAkids.com but which actually resolves (through domain forwarding) as <https://www.dss.virginia.gov/supportvakids/>.

¹⁸ For example, our desire to keep the explanations of services to essentially the same length across the three content areas on the site might arguably have short-changed services that are particularly complex or difficult explain, such as family engagement. More concretely, in Cycle 1 we accidentally left paternity services off the established list of services about which users could request information using the online contact form. This error was discovered only after the conclusion of this cycle and almost certainly affected the number of requests we received for paternity services, though by how much is unclear.

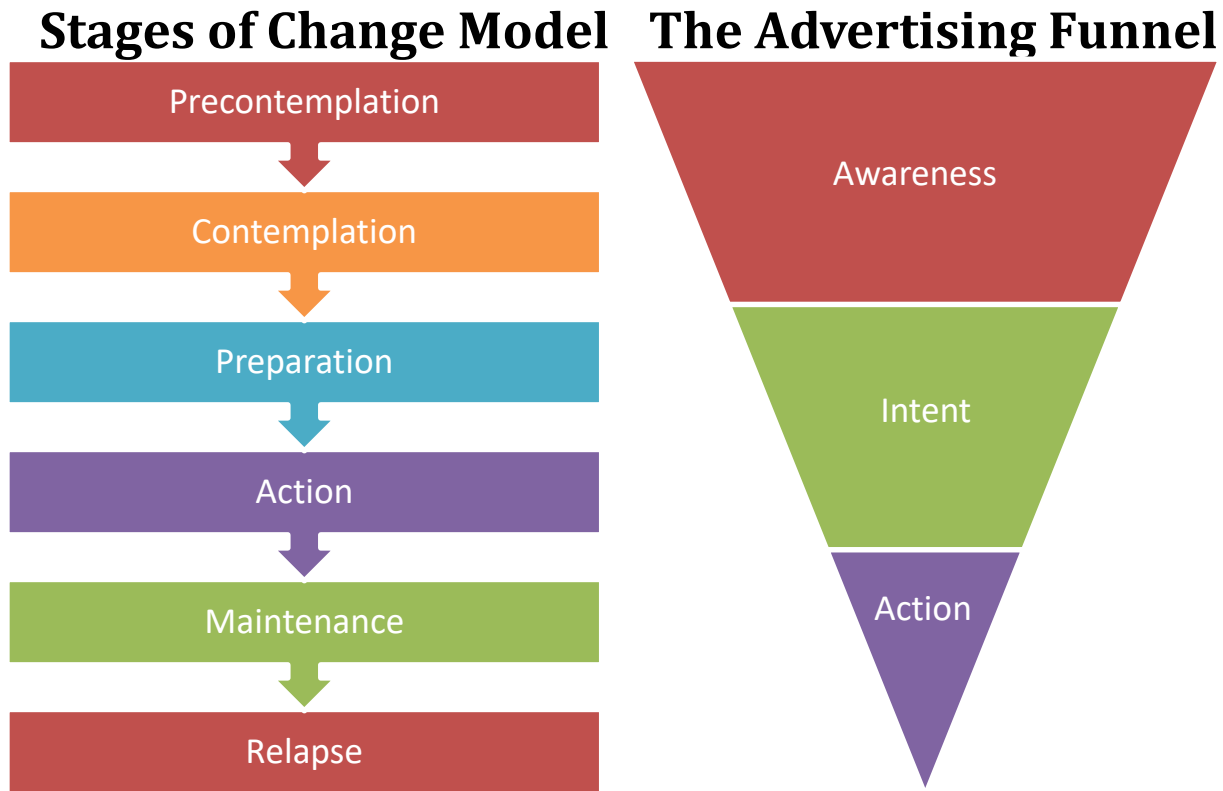


Figure 1: Comparison of the Stages of Change Model used in this project and a simplified version of the standard advertising funnel.

The organic campaign did not attempt to lead users through the funnel and instead took a broader approach to messaging, in part because it needed to lay the groundwork for a long-term approach to building its social media presence. To do that the project team and DCSE leadership agreed that, in addition to urging users to apply for child support and promoting the online child support application, the organic campaign should also highlight DCSE’s paternity and family engagement services and encourage people to follow the account and like its posts.

E. Determine Approach to Geographic Targeting for Ads in the Paid Campaign

After Cycle 1, we reviewed the available data for indications that the ads delivered during that cycle might have been more productively targeted to areas other than the 40 localities¹⁹ we originally selected. Since all of the ads had the same geographical targeting, differences in ad-performance outcomes (click-through rate, cost per click, etc.) could not be attributed to differences in targeting.

¹⁹ A complete list of the 40 localities appears in Appendix E of the Communication Plan submitted as part of this grant, along with a detailed explanation of how we arrived at this particular set of 19 cities and 21 counties across Virginia.

We then turned to examining responses to the online contact form that appears on the main project webpage. Between June 10, 2019 and September 9, 2019, we received approximately 531 submissions through that form.²⁰ Of those submissions 337 included address information that showed they fell within the targeted locations. The rest either fell outside our target areas or lacked enough data to map.

That 63 percent of the responses came from our targeted locations provides only circular evidence for the effectiveness of the targets, since finding the main project webpage either required seeing one of our location-targeted ads or receiving a direct reference to it from another user; an organic search was unlikely to reveal it. Most responses, then, would almost inevitably fall within the targets.²¹

At the same, among the group that could be mapped but fell outside our targeted areas, we saw some evidence suggesting that counties could be added based on the number of responses they generated. There were also counties in our targeted locations that generated little or no contact form submissions.

In the absence of clear data suggesting that, on the whole, the targets were not working as anticipated, the project team decided by September 1, 2019 to continue to focus on those same 40 localities for the paid campaign in Cycle 2.²²

²⁰ This figure revises the total of 554 contact form given in the interim evaluation of this project submitted after Cycle 1. Further analysis revealed that duplicates and test entries had been included in the earlier figure.

²¹ Given the location targeting and the narrow path by which most people could discover the contact form, it might seem surprising that there would be more than a handful of responses from outside the targeted localities. But the location targeting mechanism in Google (which served all of our ads in Cycle 1) does not derive from a strict measure of physical positioning, such as server or mobile phone geo-locating. A detailed explanation of ways that Google looks at location targeting is available here: <https://support.google.com/google-ads/answer/1722038>.

²² When we recently returned to this decision in preparation for Cycle 3, we again chose to target the same set of zip codes. We took this route despite some evidence from the unintended departure from our location targeting in Cycle 2 that other parts of Virginia show a high level of interest in our services. Evidence of that interest can be seen particularly clearly in the map associated with contact form responses in Cycle 2 (the third of the three maps in Appendix D). Further evidence is found in the discussion below, in section 2.V.C.IV, of the shift in geographic spending patterns that came about as a result of that accidental change in targeting in Cycle 2; the data for that discussion is broken out in Tables 1–3 of Appendix C. While both sets of data suggest that our original geographic targets did not perfectly capture the potential level of digital engagement across Virginia, we have continued to rely on those models in part because they were meant to answer a larger and more long-term question: which parts of the state have the greatest potential for growth in new child support applications, particularly among Never Assistance cases? Since the evidence we had did not clearly indicate a higher level of interest in (or follow-through associated with) applications in particular outside our target area, we plan to maintain our targeting approach for the rest of the intervention. Whether we ultimately recommend that Virginia or other jurisdictions take a similar approach after the grant will depend on whether we see any appreciable changes in the number of new cases associated with these locations.

Location targeting is not an option for organic social media campaigns on any of the platforms we used in Cycle 2.²³

Maps of contact form submissions during Cycles 1 and 2 and for the whole span between June 10, 2019 (the beginning of Cycle 1) and January 29, 2020 (the end of Cycle 2) are included as Appendix D.

F. Determining the Types of Images to Be Used in the Paid and Organic Campaigns

In Cycle 1, we tested 12 stock photos in our Google Display ads to prepare ourselves for Cycle 2.²⁴ We hoped to see results pointing us to a specific set of images or certain types of images (photos of a diverse set of children at play, for example, or a mother or father caring for their child) that appeared to increase user engagement. While certain headlines and text descriptions used in Cycle 1 were clearly more clearly associated with higher click-through rates, among other measures, the image results were less clear cut, particularly for the application content area.²⁵

We considered testing a new set of images to use in the paid campaign for Cycle 2 but felt that the potential benefits were too unclear to warrant the time required to reassess our selection criteria, generate a new set of options, and move those new images through the approval process. As a result, we decided to retest the images in Cycle 2 to see whether we might be able to draw any stronger inferences about performance.

As with the decisions around messaging, the organic side of this cycle's interventions took a broader approach to the kinds of images used in its posts. While integrating the campaign logo into its profile imagery and including photos reminiscent of the ones used in Cycle 1, the

²³ The only exceptions were the six the promoted (also called boosted) organic posts integrated into our paid campaign. For more information on those posts, see below, section 2.V.C.IV.

²⁴ Google also treated the campaign logo as an image asset that it evaluated alongside the photos but the two versions of logo that we used in Google Display ads in Cycle 1 are not counted here, though elements of them figured into all of the ads as the avatar for DCSE's social media accounts.

²⁵ We plan to test images again in Cycle 3, though this line of research is considerably more fraught than we had originally anticipated. Perhaps the greatest challenge is feeling confident that any difference in performance between two photos can be reasonably connected to the distinction we had in mind when selecting them. For example, we have tried to understand whether users tend to respond better to an image of children or of a customer service representative in an ad for applications. While simple in practice to find those types of images, the number of differences between any given example of them produces a cornucopia of confounding variables, such as the number of children or customer service people depicted in each photo, the age, race, gender, and clothing styles of those people, or even the purely photographic qualities of the images, such as the degree of contrast between the central figure or figures and the surrounding environment. This could perhaps be ameliorated through a rigid testing sequence that attempted to control and test for each of these variables but this would likely require commissioning custom photography for just this purpose—a prospect not accounted for in our budget.

organic side of this intervention drew on a wider variety of stock photos and edited them more aggressively, sometimes adding motion graphics that made them resemble short videos.²⁶

G. Determining Approach to Video in the Paid Campaign

An additional reason for testing different photo subjects in Cycle 1 (described above in section 2.IV.F.) was to see whether outcome data suggested clear paths for how we should approach the use of video in Cycle 2. Digital engagement professionals in the public and private sectors have long recognized the benefits of using video to promote content on social media,²⁷ but social media videos can take a variety of forms, including:

- Carefully shot and edited images of people directly associated with an organization or agency
- Stock footage repurposed to match a given message
- Animations that use cartoon-like characters and simple captions
- Animations that use only text and graphics
- Informal, quickly assembled video that reflects the casual nature of social media

Project staff and relevant leadership in DCSE and VDSS Public Affairs reviewed samples of various types of video and agreed that animated video appeared to offer the greatest flexibility and the clearest alignment of verbal and visual messages.

As noted above in section 2.IV.F., the organic side of this intervention occasionally posted GIFs that echoed the qualities of video, but video per se, including the animated videos created for the paid campaign, was not used in organic posts.

H. Creating Content Calendars for the Organic and Paid Campaigns

To facilitate planning and coordination between the organic campaign (developed and implemented entirely by DCSE staff) and the paid campaign (developed and implemented largely by vendor staff, with significant oversight by DCSE and VDSS), we created an online content calendar using the project management tool Asana.

The calendar allowed members of the project team to see on what day and on which platform organic and paid posts would appear and review the associated text and images. Since the organic campaign used essentially the same language and imagery on all three platforms, the calendar format provided a clear, straightforward way to review and track those posts.

²⁶ These images were animated GIFs—that is, multiple images combined and sequenced using the capacities of the Graphic Interchange Format, which can display both still and moving images as a single file.

²⁷ On Digital.gov, authors from HIV.gov, a website guided by the Federal HIV/AIDS Web Council, offer a relatively recent summary of the benefits of using video: <https://digital.gov/2017/08/07/by-the-numbers-why-video-is-effective/>.

Using an Asana calendar to track the paid campaign, on the other hand, quickly became untenable for a project that released approximately 149 ads over the course of 90 days.²⁸ Instead, the paid campaign used spreadsheets and cloud storage to bring together the information needed for each post and the schedule for when they would appear.

The paid calendar had two main components. The first listed each ad and the information needed to place it, including the textual components (headlines, descriptions, calls-to-action, links), the filename of the image or video, the start and end date of each ad, the record of which element was being tested, and the targeting information, among other details.

The second component of the paid calendar was an overall structural guide that initially broke the cycle out into two primary goals and each goal into two secondary goals, for a total four tiers corresponding to different parts of the advertising funnel.²⁹ To simplify our references to them we assigned each ad cycle a number and then a color on the calendar. Subsequent iterations of the initial cycles were indicated by adding a letter to the name of ad cycle (so the second iteration of Ad Cycle 1 became Ad Cycle 1A, and so forth).

When we introduced Lead Form ads to Facebook and Instagram at the end of December 2019 (a decision described below in section 2.V.C.IV.), we added that as a third tier within conversions. All of these elements are described in Table 1.

Table 1: Ad Cycles for Intervention Cycle 2

Primary Goal	Secondary Goal	Ad Cycle No.	Calendar Color
Branding/Awareness	Learn	1	Blue
	Like	2	Orange
Conversion	Understand Benefits	3	Green
	Apply	4	Purple
	Lead Form	5	Ochre

In the planning phase for Cycle 2, we divided the intervention into a series of four-day ad cycles,³⁰ focusing for the first sixteen days on branding and awareness and then running the conversion cycles alongside them to have one set focused primarily on new users (Ad Cycles 1 and 2) and the other on users who had already seen the earlier set (Ad Cycles 3 and 4). Carving up the calendar in this way would, we hoped, successfully bring users through the advertising

²⁸ See below, section 2.V.C.III., for a detailed discussion of the number of ads placed and the difficulties of arriving at a final count.

²⁹ We describe the advertising funnel for Cycle 2 above, in section 2.IV.D., and in our communications plan for that cycle.

³⁰ We chose four-day cycles based on Facebook’s recommendations for ad testing periods (<https://www.facebook.com/business/help/290009911394576?id=162694231107536>). For simplicity’s sake, and because all social media advertising platforms recommend frequently changing ads to keep users interested over the length of a campaign, we also planned to advertise on Twitter in four-day cycles as well.

funnel: from initial exposure to our message and brand to conversion—which in this case meant applying for child support using the new online application.

Since they were intended to appeal to two distinct audiences, we planned to have Ad Cycles 1 and 2 each run twice to build up the first audience of people new to us on social media and to allow enough time to ensure the new online child support application (the primary conversion target for this cycle) had been thoroughly tested. After we had sixteen days of user data, we planned to introduce Ad Cycle 3 to remarket to that initial audience, while running a new iteration of Ad Cycle 1 during the same four-day period. When that ended, Ad Cycles 2 and 4 would pick up where they left off.

Because it used a novel format and had a distinct conversion method as its goal, Ad Cycle 5 (our Lead Form ads) appeared at the same time as Ad Cycles 1/3 and 2/4. For similar reasons, we never planned to rotate different versions of it in four-day cycles.

In practice, we did not post four-day cycles and, because of errors in how the ads were posted, some of our conversion ads were showed to users who had not been exposed to the branding and awareness ads, thus bypassing our intended funnel process. (Both issues are described below in section 2.V.C.III.) However, we maintained the overall approach of working cyclically through two sets of branding and awareness ads while also running conversion-focused ads.

I. Registering Social Media Accounts on Chosen Platforms

We registered social media accounts on three platforms, under the following usernames:

- Facebook: @ChildSupportVirginia
- Instagram: childsupportva
- Twitter: @ChildSupportVA

We had some difficulty in registering these accounts on Facebook and Instagram and attribute those difficulties to the increased scrutiny placed on social media platforms in the wake of the 2016 election, when paid and organic social media posts had been used to disseminate misinformation.

Nonetheless, all three DCSE accounts were live by November 1, 2019, the beginning of Cycle 2.

J. Ensuring DCSE's New Online Child Support Application Is Operational

In the initial planning phases of this grant, DCSE began to investigate the feasibility of creating an online child support application in time for Cycle 2. Though not formally in the scope of work for the grant itself, DCSE's online application followed a parallel development track, much like the simplified paper application introduced around the time of Cycle 1. Project staff were kept apprised of progress, given opportunities to provide feedback and test early iterations, and

walked through the backend workflow that would turn applications into opened cases. After investigating several different implementation pathways, DCSE decided to host the online application on the same site used for its web portal (<https://mychildsupport.dss.virginia.gov/>). Images of the page and the beginning of the application are shown below as Figure 2.

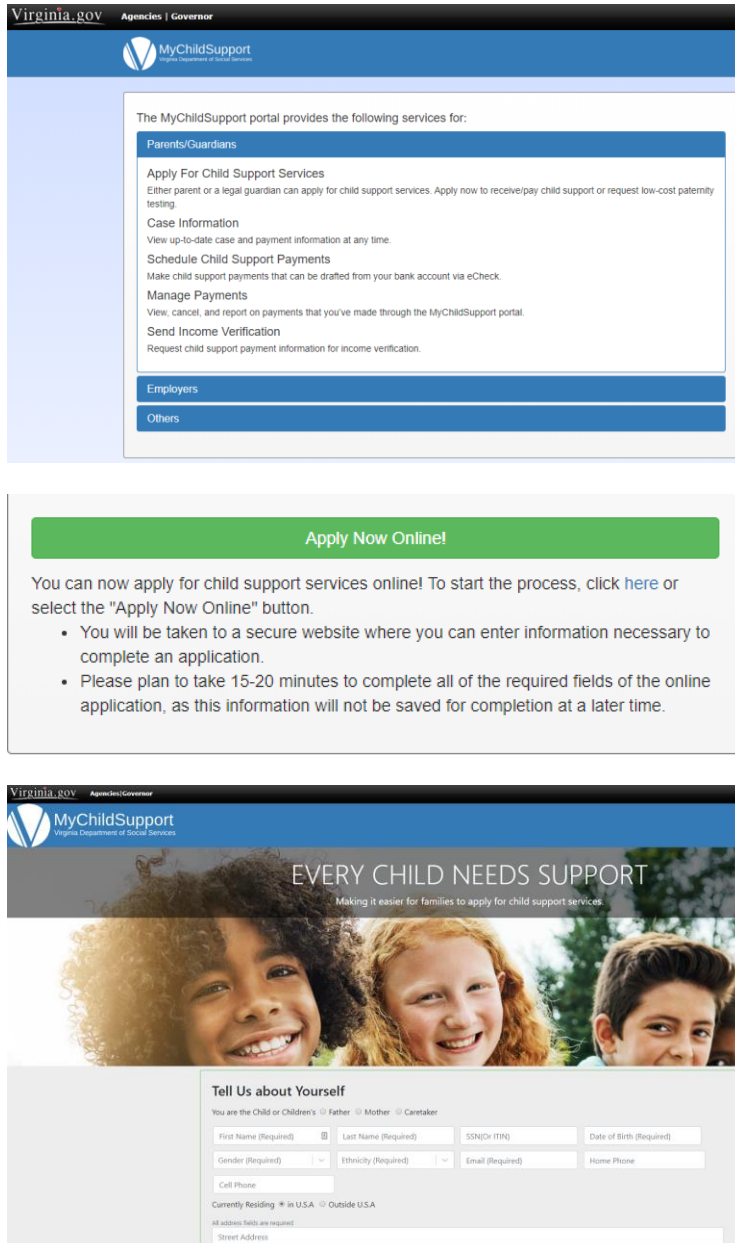


Figure 2: Top: Screenshot of part of the MyChildSupport webpage, which hosts the recently created online child support application (described under the heading "Apply For Child Support Services"). Middle: Call-to-action section ("Apply Now Online!") of the MyChildSupport page for the application. Bottom: The beginning of the application itself, which echoes the design of the supportVAkids pages.

One clear tie between the final version of the online application and this project is the header image used on the application page, which represents a slightly cropped version of the photo

on the main project webpage. The clear, approachable language used on the application form and its overall focus on simplicity also reflect a focus on user-centered design shared by the project's contact form and ad language.

To give our project a chance to test the effectiveness of our advertising, DCSE initially planned to treat Cycle 2 as essentially a soft launch period for the application, so that only users sent to the site by our ads or visitors who came to the MyChildSupport page on their own would be likely to see it. This was also intended to keep the number of electronic applications manageable while any issues with the new intake process were being addressed. Once the cycle was over and any wrinkles in the intake process were smoothed out, DCSE could then decide to promote the online application more widely (through its partners, for example, or its customer call center).

By the time Cycle 2 launched on November 1, 2019, testing of the online application was complete, allowing us to develop ads promoting it.

V. Cycle 2 Description

Cycle 2 of our project had six major components:

- 1) Creating a website to track application conversions
- 2) Developing, executing, and monitoring the organic social media campaign
- 3) Developing, executing, and monitoring the paid social media campaign
- 4) Developing, executing, and monitoring the social media hashtag campaign
- 5) Revising the budget and content calendars in response to changes in organic and paid campaigns
- 6) Monitoring and responding to social media comments

A. Creating a Website to Track Application Conversions

From a project perspective, the new online child support application created one significant hurdle: because it was hosted on the MyChildSupport site (mychildsupport.dss.virginia.gov) rather than as part of the DCSE family of pages (branching off of dss.virginia.gov/family/dcse), we could not easily add to it the same tracking code used on the main project webpage. Without that code, we would have no mechanism for recording users who visited the page or converted by submitting an application—both critical elements in our larger marketing plan.³¹

The issue here was one of user security rather than technical limitations. Because the MyChildSupport site primarily functions as a client portal, there were concerns about adding a third-party tracking tool, especially one necessarily administered by another third party (the

³¹ During Cycle 1, we came believe that conversion tracking was an essential ingredient for any online advertising campaign. For a fuller description see subsection 5 of section 3.III. of our interim evaluation for Cycle 1.

project’s vendor), to a site where clients enter usernames and passwords and view or enter personally identifiable information.³²

As a compromise, the project team decided to create a new page on the main project site that would provide additional information about the application process, reinforce our brand identity (including our social media presence), and then link to the online application. Clicks on the link made through the button marked “Apply Now!” could then be recorded as conversions by our tracking tool. Creating a page specifically devoted to explaining the new online application would also give us a place to post information about the application process, including narrative text or videos explaining how to apply, what paperwork is needed, and so forth.

Since we needed to tie the new page directly to the larger project, we planned to show the address for it as [apply.supportvakids.com](https://www.dss.virginia.gov/supportvakids/apply/) in our ads. Users who clicked on that link or typed it in, however, would be redirected to <https://www.dss.virginia.gov/supportvakids/apply/>.

To simplify the development and review process for the new page and underline its ties to the main project webpage, we proposed a design that carried over the layout, imagery, color scheme, and much of its text from the existing project webpage. We gave it the headline “Applying for Child Support in Virginia Has Never Been Simpler” so that there would be a clear tie to the overall theme of simplicity.³³

The VDSS Public Affairs unit built the page in a matter of hours. The project team then tested the tag used to record clicks on the “Apply Now!” button. The project’s logo and the social media information for DCSE were then added to the bottom of the main project webpage and this new page—referred to here as the Apply page.

All of this was completed before Cycle 2 launched on November 1, 2019.

³² To be clear, the tracking tool used by project team, Google Tag Manager, does not collect login information or personal data, only aggregated information about certain user behaviors. The concerns about MyChildSupport were raised out of a reasonable sense that the sensitivity of child support information requires programs be extremely vigilant about security, and project staff felt that, given the straightforward workaround available and their shared sense of the importance of user privacy, there was no reason to press the case for adding the tracking code. For details about the privacy policies associated with Google Tag Manager and its adherence to the third-party certification standards, see <https://support.google.com/tagmanager/answer/9323295?hl=en>. For an overview of Google Tag Manager, see <https://support.google.com/tagmanager/answer/6102821?hl=en>. We should also note that toward the middle of the cycle, once Virginia’s security staff had a chance to fully review the question of adding conversion tracking to that site, they approved it for use. It may be implemented on the MyChildSupport page for Cycle 3 but was not in Cycle 2.

³³ The headline also put *Virginia* in close proximity to *Child Support* in recognition of the importance of *Virginia Child Support* as a keyword phrase in search engine optimization information, as we described in our interim evaluation of Cycle 1.

An image of the Apply page appears below as Figure 3. The text of the webpage appears as Appendix A.



Figure 3: Screenshot of the new Apply page that branches off the main project webpage and links, through the "Apply Now!" button, to the online child support application.

B. Developing, Executing, and Monitoring the Organic Social Media Campaign

I. Developing Initial Posts and Determining Review Process for Organic Campaign

Having established the content areas that the organic campaign would address, we began to develop proposed text and images that would be combined for posts on Facebook, Instagram, and Twitter.

Initially DCSE project staff created a set of sample posts for review by leadership within their own program and staff and leaders in VDSS's Public Affairs office. One of the goals, especially during the initial period of the campaign, was to establish a consistent and engaging voice for the organic posts, one that reflected the overall approach of its parent agency but distinctive enough to encourage its own following. Understood in a holistic sense on social media to include the kinds of posts shared and images used, a consistent voice can "build enduring relationships with existing followers along with attracting a new audience. People feel comfortable when they know what to expect and can relate to content that's consistent in its delivery."³⁴

³⁴ Giovanbattista Cimmino (2019). "The Importance of a Strong Brand Voice on Social Media (and How to Develop One)." *Business Matters*. <https://www.bmmagazine.co.uk/columns/the-importance-of-a-strong-brand-voice-on-social-media-and-how-to-develop-one/>.

As the organic campaign continued, DCSE took over more and more of the content development process, with VDSS Public Affairs limiting its role largely to review and approval of what child support staff had created. Meanwhile, project vendor Grays Peak Strategies also provided a modest amount of feedback on these posts and the structure of the campaign generally.

This division of labor—leaving the creation of organic content largely to VDSS staff and letting Grays Peak Strategy lead the development of ads—reflected both the skills of the various staff involved while also testing DCSE’s ability to sustain an organic presence on these platforms.

The paid campaign also paid to promote six organic posts. While these can be considered advertisements, they did not pass through the usual ad development process and were instead selected by the project team based on theme and performance. More details on this process appears below in section 2.V.C.IV.

II. Launching, Monitoring, and Maintaining the Organic Campaign

The first organic posts to DCSE’s social media accounts on Facebook, Instagram, and Twitter appeared on the morning of November 1, 2019. Images of those posts appear below in Figure 4.

Those first posts were accompanied by direct outreach to partners (such as other state and local government agencies, nonprofits, and other child support jurisdictions with active social media accounts) to encourage them to like the post and to follow the account. We had already identified many of these partners before the cycle began by generating lists of other child support programs with a social media presence, soliciting input from other parts of DCSE about agencies or organizations important to the program, and identifying child support leaders and potential public and nonprofit sector partners. We followed as many of them as possible on each of the platforms and, once the cycle began, actively liked, shared, and retweeted their content as part of our broader engagement strategy.



Figure 4: The first organic posts to DCSE's accounts on Twitter (top), Facebook (bottom left), and Instagram (bottom right).

Organic posting continued across all three platforms, generally at twice-weekly intervals, throughout the cycle. We worked directly within the platforms rather than using a social media tool, such as Hootsuite, so we could be sure the content was suited to the individual site, but took advantage of planned posting capabilities within the platforms, when available, to better control the timing. To make the posts visually appealing we integrated stock photos and other graphics into the messaging, as noted above in section 2.IV.F.

In the end the organic campaign made 24 posts to each of the three platforms between November 1, 2019, and January 29, 2020, for a total of 72 over the course of the 90-day cycle. We hewed to our rough schedule of approximately two posts per week (1.8 per week for 13 weeks) or approximately one post every four days. These organic posts can be divided into six categories:

1. **Account Promotion** (encouraging users to like or share the post or follow the account)
2. **Applications** (encouraging user to apply, usually with a reference to the online application)

3. **Family Engagement** (informing users of DCSE programs connecting them to employment assistance and other nontraditional child support services)
4. **Paternity** (informing users about the benefits of paternity establishment)
5. **Payments** (informing users of the variety of payment options available)
6. **Holiday Notices/Customer Service Updates** (informing users of temporary changes to the hours of public-facing locations, typically beginning with holiday-related greetings)

A summary of posts by platform and theme is in Table 2.

Table 2: Organic Campaign Posts by Platform and Theme

Theme of Organic Post	Number of Posts			
	Facebook	Instagram	Twitter	Total
Account Promotion	4	4	4	12
Applications	7	7	7	21
Family Engagement	2	2	2	6
Paternity	5	5	5	15
Payments	2	2	2	6
Holidays/Updates	4	4	4	12
Total	24	24	24	72

Six of our seventy-two organic posts (four from Twitter, two from Facebook) received additional exposure through promotions as part of the paid campaign. We selected the six posts because they both supported the paid campaign’s focus on applications and had already demonstrated relatively high organic engagement rates.

Additional data on organic posts appears below, in section 3.III.A.

Information about monitoring and responding to comments—a crucial element of monitoring the organic campaign—is found below, in sections 2.V.F. and 3.III.B.

C. Developing, Executing, and Monitoring the Paid Social Media Campaign

I. Developing Initial Ads for Review by Project Staff and Leadership in DCSE and VDSS Public Affairs

At weekly or biweekly meetings of the core project staff, we considered general approaches to messaging for paid content and previewed potential still images and videos. Grays Peak Strategies staff then proposed a specific combination of text (for headlines, descriptions, and calls to action) and photos or videos for review and submitted those to the project team through our project management tool. Once the cycle began, our ad development process

was also informed by performance data (click-through rates, costs per click, conversion rates, etc.) from the various platforms.

As noted above in section 2.IV.D., we focused the paid campaign on encouraging new child support applications, as we had in part done in Cycle 1 but with a new emphasis on DCSE's online child support application. There were essentially three components of the ads created to convey those messages:

1. **Images.** As explained above in section 2.IV.F., we chose to test the same set of still photos we used in Cycle 1. Each Ad Cycle featured at least one set of ads using still images.
2. **Videos.** We tested two messaging streams against each other in our videos:
 - a. **Child Support Made Simpler.** This set of videos emphasized the online application and the support services provided (such as customer service agents for parents applying for services or resources to help noncustodial parents find work). They also used behaviorally informed statements (such as "Kids' little expenses add up" or "Your child's needs can't wait!") to motivate users to follow the call the action.
 - b. **Child Support Facts.** While these generally included a direct reference to applications in its primary message or call-to-action, this series of videos took a broader approach to brand awareness, hoping to frame the child support program as a source of neutral, purely factual information. We also hoped users who found the information surprising would be more likely to share the videos. We considered taking the myth-busting approach used in other jurisdictions but felt that this essentially negative frame seemed ill-suited to social media, where research shows users are more likely to share positive messages.³⁵

The two series of videos were further subdivided by where they fell on the advertising funnel. Videos intended for branding and awareness (Ad Cycles 1 or 2) were typically 9 seconds or less and usually made only one claim, such as "NEW online application" or "Child support orders always include medical support" or "Fact. Child support: Now you can apply online!" Videos associated with conversions (Ad Cycles 3 or 4) ran between 12 and 16 seconds and included several points, at least one of which we hoped would be familiar to viewers from the earlier, shorter version.

The distinctions between the two series in terms of content and framing masked their overall similarity. All of the videos produced for this intervention:

³⁵ See, for example, Jonah Berger and Katherine L. Milman (2011). "What Makes Online Content Viral?" *Journal of Marketing Research*. <https://bit.ly/PositiveSharing>.

- Were created in both horizontal and vertical formats, so they could be delivered equally well to desktops and mobile devices
- Used animated figures and text created using the online animation tool Powtoon
- Featured upbeat music in the background, with no voiceover or other essential sound³⁶
- Began with a brief title sequence (lasting 1 to 3 seconds)
- Included a small, static image of the project logo at the bottom of the frame from the title sequence through the end sequence (termed an outro). Horizontal videos used a horizontally oriented version of the log in the lower right of the frame, while the logo fell in the center of vertical videos. We generally used the versions of the logo featuring the name, rather than those that included both the name and the image of the penguins³⁷
- Concluded with an outro (lasting at least 2 seconds) that gave the project website URL (supportVAkids.com) and logo, and showed the call-to-action corresponding to the overall call-to-action (e.g., an ad that used “Learn More” as its call-to-action was accompanied by a video with “Learn more at” in its outro)

3. **Text.** Initially the language we used in the videos and in the headlines, descriptions, and calls to action in the ads came directly from the text that appeared in our Search or Display ads.³⁸ Indeed, the few deviations from that earlier language that we made were done to align the text of the ads with the titles and captions in the videos. For example, the “Child Support Made Simpler” ads began with the phrase “Moms and Dads” because we wanted to see whether it might improve salience for users who might glance at the video before looking at the text. We also hoped to activate users’ view of themselves as parents. The “Child Support Facts” series, meanwhile, used the word “Fact” prominently in its title section and on each scene describing a fact.

³⁶ We took this approach in part under the assumption that social media videos are largely viewed without sound. Articles on best practices for social media video often make this point, which aligns with our team’s personal experience. But we should acknowledge that the data in support of it appears to be limited to a single article that attributes its information to “multiple publishers.” This apparently means the websites Mic and PopSugar and the advertising agency MEC North America. Sahil Patel (2016). “85 Percent of Facebook Video Is Watched without Sound.” *Digiday*. <https://digiday.com/media/silent-world-facebook-video/>.

³⁷ A sampling of the various versions of our project logo is given as appendix G in our interim evaluation for Cycle 1.

³⁸ The language used in the ads from Cycle 1 drew on a variety of techniques associated with behavioral economics, a topic we discussed at length in our interim evaluation for Cycle 1. Since our Cycle 2 ads used essentially the same language and overall messaging principles employed in Cycle 1, we inevitably drew on the same behavioral techniques. We have therefore omitted a detailed description of our behavioral approaches from this report, despite their central place in our ad development process.

II. Launching the Paid Social Media Campaign

Our paid campaign began largely as planned on November 1, 2019 and continued through January 29, 2020. The execution of the campaign was successful by that measure and others, not least in possibly contributing to a slight reversal of trends in child support applications, as we discuss below in section 3.III.D. Images of two representative ads appear below as Figure 5.

Yet the execution of the paid campaign also hit a number of significant bumps along the way.



Figure 5: Left: Mobile version of Facebook and Instagram ad from early in Cycle 2 using a still image. Right: Desktop version of the same ad, showing a frame from one of the animated videos.

III. Delayed Approvals for Ads on Facebook and Instagram

The first problem became clear on November 1, when we realized that there were significant delays in receiving ad approvals from Facebook and Instagram, particularly for video ads.³⁹

The experience of Grays Peak Strategies had, up to this point, been that ads submitted to Facebook, Instagram, and Twitter received only nominal review before being approved at most a few hours later. That was true regardless of whether the ad used a still photo or a video, and for Twitter, that continued to hold true during Cycle 2.⁴⁰ However, similar to the problems we faced when first trying to establish a DCSE account on its platform, Facebook held all six of our initial ads for review for more than 24 hours without giving us any opportunity to discuss the situation directly with a customer service representative.

Instead, we received an automated email about 24 hours after the ads were placed saying that our video ads had violated their policies against implying that Facebook had endorsed our message. Neither the text of the ad nor the captions in the video actually suggested that, and the email provided no specific objections. We eventually decided that they might have

been objecting to our use of the Facebook and Instagram logos at the end of the video. We had considered the use of these and the Twitter logo to be a way of reinforcing our organic presence on the various platforms, but we immediately removed them from the videos and resubmitted the ads. Though we could never be sure that was the issue, we never used those logos again and never ran up against that same objection.

The first approved Facebook ad used a still photo that we also used on the main project website—a sign, we felt, that they remained somewhat wary of us as a new account and looking for ties between our ads and our website that suggested we were a legitimate program and advertiser. The other ad with a still photo was approved the next day, November 2, along with the first approved video ad: a horizontally oriented one from our Child Support Facts series, described above in section 2.V.C.I. Two more days passed before Facebook approved the other three videos ads, with no information from the company in the meantime or afterward and no meaningful opportunities to appeal their decisions or receive the feedback needed to avoid these problems going forward. This meant that half of the posts in this first ad cycle were shown on the last day instead of spread out evenly over several days.

These kinds of delays—and accompanying silence—continued for almost the rest of the cycle with video ads posted to Facebook and Instagram. To address the delays, the campaign first extended Ad Cycle 1 from four days to eleven so that we would have enough time to get the next set reviewed, approved, and submitted well in advance. This proved somewhat successful for Ad Cycle 2, which was intended to last only four days. The content for Ad Cycle 3, however, was not developed, approved, and submitted quickly enough, and only still photo ads were approved and visible to users on the first two days of that cycle.

Despite deciding to move away from our original intention of beginning a new ad cycle every four days, similar problems recurred for the first 60 days of the campaign:

- On November 21 and 22, when we attempted to simply extend the existing ad set for Ad Cycle 3 but were pushed once again into the review process
- On December 8 and 9, when Ad Cycle 1A and 3A were running concurrently
- On December 19, when transitioning to Ad Cycles 2A and 4A
- On January 2, when transitioning to Ad Cycles 1B and 3B

³⁹ Facebook owns Instagram so ads for those platforms pass through the same ad development and management platform and presumably the same review process. In this report we therefore tend to refer only to Facebook when discussing advertising policies and processes that relate to both platforms.

⁴⁰ Toward the end of Cycle 2, however, the paid campaign began to hit an even more intransigent and difficult to understand problems with video ads on Twitter, as noted later in this same section.

The only way around this problem was to submit a new set of ads about five days in advance, which we succeeded in doing for the few remaining ad cycles.

IV. Monitoring the Paid Social Media Campaign and Adjusting Ads and Placements

The project team met weekly and sometimes biweekly to discuss all aspects of the intervention, but with particular emphasis on the paid campaign. We used this time to discuss key performance metrics (such as impressions, clicks, cost per click, click-through rates, conversion rates)⁴¹ and strategized ways to improve outcomes. All along, we were also comparing those performance metrics to our goals.

Our first set of changes came with the language used in the ads we released in the “Child Support Facts” series that started on December 2, 2019. Without altering the substance of the underlying content, the project team worked with leadership in the VDSS Public Affairs office to simplify the language and bring the syntax of each of the facts more clearly into parallel. The ads with this new language debuted on December 2 in Ad Cycles 1A and 3A.

We made only a small number of additional changes to the language of the ads over the remainder of the cycle, all of them minor matters of style or punctuation.

One area of particular concern throughout the cycle was ad placement, both across and within the three platforms. For the first month we let Facebook’s algorithm decide where to place our ads within its own platform and Instagram. This approach granted them the power to determine whether a given ad should appear on Facebook or Instagram, or within Facebook on its right column (a desktop-only option) or in a news feed for mobile phones—to list only a few among the 19 distinct placements our ads eventually received.⁴² It also led to our ads appearing in two places that seemed out of place for our approach to this intervention: Facebook Messenger (a messaging app) and the Facebook Audience Network (essentially a set of contracted display advertising opportunities on apps and websites in keeping with the model of the Google Display network). We excluded these placements after the first week.

Thirty days into the intervention we reviewed the data in more detail to assess the effectiveness of letting Facebook determine ad placements. One conclusion was that our ads were being delivered quite unevenly between Facebook and Instagram. Between November 1 and November 30, 2019, our ads received only 8,503 impressions on Instagram compared to 299,042 on Facebook—a difference of 3,500 percent. We decided to change our settings

⁴¹ For a glossary of common online advertising terminology, see <https://www.iab.com/insights/glossary-of-terminology/>.

⁴² Facebook provides an interactive overview of the placement options: <https://www.facebook.com/business/ads-guide>. Seeing all of the options, though, requires using the menu at the top of the page and selecting the ad type (video, image, etc.) and then reviewing further options under that page’s “Choose your ad placement” menu.

going forward to give each platform its own budget, rather than letting Facebook choose how to allocate a single campaign ad spend. This succeeded in bringing the number of impressions across the two platforms into rough parity, but we never saw any positive changes to click-through rates, conversions, or other performance measures.

Our review of program data at the 30-day mark also ultimately led to another change in placement, this time concerning Twitter. During November 2019, our ads received 167,880 impressions on Twitter but only 235 clicks on the link to the website, for a click-through rate (clicks divided by impressions) of just 0.14 percent—only a fraction of our benchmark at that time of 0.84 percent and our goal of 0.93 percent. Our ads on Facebook and Instagram, by contrast, were roughly in line with our goals and benchmarks and it seemed possible to close the gaps that existed.

After roughly two more weeks of advertising without seeing any improvements or understanding why the ads fell so far from the mark, we decided to suspend advertising on Twitter altogether on December 12, 2019, and focus on Facebook and Instagram. We eventually returned to Twitter advertising on January 10, 2020, for reasons described below.

During the second month of the campaign, we also evaluated the effectiveness of the six Twitter and Facebook posts that had originally been posted as part of the organic campaign but which we had promoted using funds from the paid campaign. While paying to promote them inevitably led them to be seen by a larger number of users than their organic peers, they did not lead more users to follow the account, and their performance in other respects was generally below that of posts created purely as advertisements.⁴³ Based on an initial analysis of these results, the project team decided to no longer promote organic posts as part of the paid campaign.

⁴³ The organic posts we promoted on Twitter received a click-through rate of 0.11 percent, somewhat lower than average click-through rate of 0.13 percent received by our regular paid posts during that first period. At the same time, promoted tweets received a markedly higher overall rate of engagement (a Twitter-specific metric that includes different types of clicks, as well as profile views and other measures). The average engagement rate for promoted organic content on Twitter was 2.5 percent, compared to the average engagement rate of 1.4 percent for our paid posts during that first period on Twitter—a difference of 76 percent. The organic posts that we promoted on Facebook fared less well. While our click-through rate on Facebook, Instagram, and related properties was 0.76 percent for posts developed purely as paid advertisements, our two promoted posts received a strict click-through rate of 0.19 percent. The results of the broad click-through rate that Facebook emphasizes—a metric roughly equivalent to Twitter’s engagement rate—were somewhat less lopsided, with the promoted organic posts receiving a broad click-through rate of 0.84 percent versus 1.33 percent our regular paid post. Results from such a small set of promoted posts are hardly definitive, and the project team has discussed possibly testing them again in Cycle 3, depending on the performance of our regular paid social media posts during that cycle. If we do, we plan to do so over a longer period, so we can develop a more rounded sense of the potential value of this method.

Toward the middle of December 2019, largely to address the relatively small number of conversions we had seen up to that point, we decided to experiment with a different ad format on Facebook and Instagram. Called a Lead Form ad, this format appears at first like a regular image- or video-based ad but instead of leading users out of the app to a website that asks them to complete a contact form, it includes a contact form in the ad. One advantage of placing this type of ad in Facebook is that it allows mobile users to stay inside the mobile app rather than exit to a browser, thus building on the path dependency created by the app. It also allows for user information to be pre-populated, saving them time and potentially increasing the accuracy of the information. All of these advantages should, in theory, smooth the path to conversion—which in this case is the user submitting the form.

One disadvantage, however, is that Facebook stores the responses to Lead Form data within its own platform. This means that data from these forms has to be retrieved separately from the data from the contact form. Facebook also deletes the data after 90 days, a measure that provides an element of privacy protection after time but also puts an additional burden on the program to store and secure the data itself.

The Lead Form ads closely followed the language and structure of the contact form we developed in Cycle 1 and which users completed in this cycle as well. An image of a draft Lead Form ad appears below as Figure 6.

Virginia Division of Child Support
Calling all Moms and Dads!

- Child support orders always include medical support
- No fee for new applicants
- Child support is linked to kids doing better in school

Complete the form and we'll respond within two business days.

Full name
James Robert

Email
jrobert@test.com

Phone number
US +1 (765) 555-1212

What can we do for you?
Help with an existing child support case

Done

Figure 6: Sample of a Lead Form ad, using pre-populated information and a menu for users to answer questions.

We ran Lead Form ads from December 28, 2019 to the end of Cycle 2. During that time, the ads generated only 37 leads, at an average cost of \$46.35 per lead. While higher than the

\$13.44 paid per conversion for Google Search ads in Cycle 1, the cost per conversion for our Lead Form ads was still slightly below the overall cost per conversion of \$66.89 for all ads across Cycle 2 (a figure that combines results from Facebook, Instagram, and Twitter but excludes the Lead Form ads themselves). It was also far lower than the \$546.52 per conversion for Google Display ads in Cycle 1. The average click-through rates for our Lead Form ads was 0.95 percent across Facebook and Instagram—or 28 percent higher than our overall click-through rate of 0.74 percent across those platforms.

Sixty days into the cycle, after another in-depth review of performance data, we made a second set of changes to placements within Facebook and Instagram. After reviewing placements that generated a large number of impressions but a relatively small number of clicks to our website, resulting in a low click-through rate, we decided to exclude Facebook’s right column placement and Instagram’s Explore area from our ads.

During the last month of the campaign we also discovered a problem with the demographic and geographic targets that had been affecting almost all of our ads from the beginning of the cycle.

The problems with demographic targeting were primarily related to Twitter, where we realized that a staff member from Grays Peak Strategies had decided to narrow the audience by adding what he believed to be relevant demographic criteria to the ad targeting. While we had planned to introduce no demographic targeting in these ads—the same approach we had used in Cycle 1—the vendor’s staff member decided to exclude users that the platform categorized, based on its own data, as married and include those who fit these additional targets:

- Moms of preschool kids
- Working-class families
- Working-class moms
- New parent 10, 11 or 12 months
- Single parent
- Moms of high school kids
- New parents
- New parents 6 months or less
- New moms
- New parent 7, 8 or 9 months
- Moms of grade school kids⁴⁴

⁴⁴ The language used in this list comes directly from Twitter’s own demographic categorization. Additional information about Twitter’s ad targeting options is available at <https://business.twitter.com/en/help/campaign-setup/campaign-targeting/geo-gender-and-language-targeting.html>.

This significantly changed our targets and likely led to poorer performance on that platform.

The changes to our location targets were equally significant. As we learned around January 6, 2020, almost all of the ads posted on Facebook, Instagram, and Twitter up to that point had been delivered to users across all of Virginia, rather than in the 40 targeted localities used in Cycle 1 and noted above in section 2.IV.E. The only exceptions were on Facebook and Instagram in Ad Cycle 1, which followed the prescribed geographical targeting.

We moved immediately to correct these errors on Facebook and Instagram and decided to return to advertising on Twitter in the next ad cycle. On January 10, 2020, we were again advertising across all three platforms and adhering to our intended targeting options.

We also looked once again at the performance data to see whether we could come to any conclusions about how these problems had affected the cycle's outcomes to that point. Based on that analysis we came to speculate that these problems might have led to a lower click-through rate on most of our ads, since the platforms' algorithms were serving an audience located outside of high-need areas. That lower level of salience might reduce the likelihood that users would click on the ads. The demographic settings on Twitter, meanwhile, skewed the targeting in less predictable ways by simultaneously excluding relevant users, such as married people who are eligible for child support, and focusing too narrowly and somewhat arbitrarily on a set of the users that remained.

Our suppositions about the effects of the altered demographic targets on Twitter could not be confirmed, since we had no demographic data for comparison. However, our ideas about the geographic targeting could be tested to some degree against comparative data from Cycle 1.

To do that we looked at ad spending by Designated Market Area (DMA) in Cycle 1 versus the same time point in Cycle 2. We have provided part of that data in Appendix C, which includes three tables ranking the ten most important DMAs by the amount spent during the two periods.

The differences in ad spending patterns between Cycle 1 and Cycle 2 are particularly marked on Twitter, where we spent 38 percent of our ad budget on the Washington, DC market area alone. During that same period we had spent only 16 percent of our ad budget on Facebook and Instagram in that same market, and only 6 percent Cycle 1. This suggests that the statewide geographic targeting led the algorithms to serve more ads in the Washington, DC market than it would have using the correct targets. This, in turn, would seem to have made our spending on Twitter less efficient than it might have been, even though the total amounts would presumably have been the same.

While significant and almost certainly a contributing factor to the low click-through rates we saw with our Twitter ads, these problems did give us some evidence that the geographical targets we used in Cycle 1 and thought we carried over into Cycle 2 might be improving the performance of our ads, as noted below in this same section.

The newly corrected ad targets came to the ongoing ad placements on Facebook and Instagram on January 10, without any clear changes in performance. When we resumed advertising on Twitter that same day, however, we immediately saw a marked improvement in performance. Under the old targeting model, the highest click-through rate of any Twitter ad up to that point was 0.21 percent; under the newly corrected targeting, we immediately began to see click-through rates of nearly 10 percent.

At first we attributed this change as evidence that our original demographic and geographical targeting were indeed more effective than the ones that we had used initially. Yet after a few days of seeing similar results, we began to wonder whether the change was too dramatic to be a reflection of real user engagement, particularly since the only change was to impressions and clicks (and hence the click-through rate), not to any other of our performance metrics that we were investigating, particularly conversion rate.

To determine whether these clicks might reflect real user behavior, rather than an error on the part of Twitter or the deliberate workings of a program designed to exploit ad clicks,⁴⁵ we compared the number of clicks that these ads said they were sending with the ones recorded by Google Analytics on the destination webpage. While not perfectly in line, the overall volume was roughly the same, leading us to conclude that the performance boost was grounded in the behavior of real users.

This conclusion seemed particularly likely given our results during the first part of Cycle 2, which showed no apparent hints of click fraud or other malicious behavior. Also, we could not discern any motivation for click fraud if the clicks were all being made on Twitter's own platform: Twitter would be the only entity to benefit from advertising expenses, so no third-party site would benefit, as is typically the case with ad-related click fraud.⁴⁶

⁴⁵ These types of programs are often referred to as click bots. A blog post from the computer services provider CloudFlare offers clear overview of click fraud and click bots: CloudFlare (n.d.). "What Is Click Fraud? | How Click Bots Work." <https://www.cloudflare.com/learning/bots/what-is-click-fraud/>

⁴⁶ In fact, as noted below in section 3.III.C., the clicks all came from Twitter's third-party advertising distribution network called the Twitter Audience Platform. Unfortunately, we realized this only after the conclusion of the cycle.

As we began to analyze the results of the intervention, however, we began to realize that the patterns we saw were consistent with click fraud. See section 3.III.C. for an explanation of what drove us to this conclusion.

Overall, the paid campaign proved challenging, both because of the problems described above and because of the volume of ads placed in a relatively short time. According to Twitter’s ad management tool, we posted 50 ads during Cycle 2. The analogous tool for Facebook and Instagram lists 99 ads associated with those platforms. On the other hand, the spreadsheet we used as a content calendar for our paid content records 78 ads during the same period, each of them with a split (or A/B) testing option.⁴⁷ The difference derives largely from the difference between recording each “A” and “B” ad as a distinct entity (149) or seeing them as a single ad that comes in two versions (78).⁴⁸

In hindsight, this high ad volume likely led to many of the problems described earlier, if only because of the complications of tracking the many details associated with each ad across multiple platforms. This point is fleshed out to some degree below, in section 2.V.E.II, which describes the modifications made to the content calendar.

By the end of Cycle 2 we also created roughly 60 short videos for integration into our ads. Two representative examples include:

- A vertical video from the “Child Support Facts” series that was used in an iteration of Ad Cycle 2 (a branding and awareness cycle): <https://youtu.be/YbZCI97Dznw>
- A horizontal video from the “Child Support Made Simpler” series that was used in an iteration of Ad Cycle 3 (a conversion cycle): <https://youtu.be/E4WRSc-JPZg>

D. Developing and Executing a Social Media Hashtag Campaign

From the initial stages of planning Cycle 2 we intended to develop a social media hashtag⁴⁹ to encourage other users to share positive content related to child support in Virginia. After some discussion, we settled on #supportVAkids, which both reflected the overall campaign

⁴⁷ In our interim evaluation for Cycle 1, we had included tables listing the text and other ad assets used to build our ads. Because of the far higher number of ads in Cycle 2, we chose not to do that in this evaluation, since the table would run to dozens of pages and would not, in any case, include the videos that we saw as central to this campaign. Readers interested in this information can request it through the key contact person listed at the head of this report.

⁴⁸ The count of the number of ads placed becomes exponentially larger if we treat each of the 19 distinct placements across Facebook and Instagram as its own ad. Using this approach, the 99 ads we tallied through those platforms suddenly becomes 1,881—a number that would be lowered somewhat by our eventual reduction in the number of placements.

⁴⁹ Used across many social media platforms, hashtags are words or phrases that allow users to group otherwise unrelated content across that platform. It typically begins with a pound or hash sign and does not include spaces.

name and suggested a broadly appealing goal that might get users to think beyond any preconceived ideas about the child support program.

Virtually every post we made across both the organic and paid campaigns used that hashtag.

E. Revising the Budget and Content Calendars in Response to Changes in Organic and Paid Campaigns

I. Revising the Budget

Once the campaigns were underway, we periodically entered actual spending amounts into the budgeting tool⁵⁰ as a final mechanism for tracking our spending.

Overall, the tool met our need for a simple mechanism for planning and tracking our ad spending. At the same time, setting the budget on a weekly schedule fairly early in our planning clashed to some degree with later plans to swap in new ad sets, with new budgets, on days that often did not correspond to the beginning of a week. This led to confusion about budgeting during some ad cycles and contributed to our overall tendency to under-spend in this cycle relative to our overall targets.

Our original budget projected \$33,050 for ads in this cycle. We reduced that to \$29,060 before the start of Cycle 2, and in the end spent \$21,984.83, for a difference of \$11,065.17 below the original budget projections.

We plan to apply the remainder to Cycle 3.

II. Revising the Content Calendars

Both the paid and social media content calendars received regular updates during the intervention. The organic campaign managed to use the original Asana calendar successfully across the cycle's full 90 days. The paid campaign continued to use the spreadsheet we initially created, but over time it became clear that our tracking tools were not performing as well as we had hoped.

F. Responding to Comments on Posts in the Organic and Paid Campaigns

How to address negative social media comments was among the most important topics we discussed with the social media and communications advisors (noted above in section 2.IV.B.) that we consulted when planning this cycle.

⁵⁰ Since most online ad platforms use a performance-based cost model, where ad costs are determined either by user action (clicking on ad, for example) or by the platform's ability to find an appropriate user to whom the ad can be displayed, the actual ad costs often fall below the budgeted amount. This is not true of traditional print or broadcast advertising, which give fixed prices based on size (as in newspapers) or time (as with radio and television), among other more or less definite, clearly predictable criteria.

We received more or less the same general guidance from all of them:

- **Not many people make negative comments.** This was usually attributed to posting positive content and thus not opening up any clear path for disagreement.
- **Responding to negative comments quickly and visibly is the best way to handle it.** The goal was not to address their case-specific issues online but to tell them to send contact details by direct message. A customer service representative would then follow up by phone. This approach would show that the program is responsive, while also addressing the commenter’s specific issues with all of the necessary privacy protections.
- **A small number of users—perhaps only one or two over the lifetime of the account—will respond negatively to many or even all of your posts and there is little if anything you can do.**

Our experience in Cycle 2 confirmed some elements of these claims. Relative to the number of impressions we received or the overall reach of our paid and organic posts in Cycle 2, we did indeed receive only a small number of comments (as in the first point above) and of those, an even smaller number consistently left negative comments (as in the third point). At the same time, we remain unsure that leaving visible responses to negative comments (as recommended in the second point) clearly improved any individual user’s experience or provided any broader benefits in the program’s engagement with program participants, largely because these benefits remain part of the affective or cognitive kind of engagement that was essentially invisible to us in this intervention.⁵¹

Of the 74 combined organic and paid posts we made on Twitter, only a handful received comments and most of them were negative. Comments on Instagram were similarly sparse, but they were far more common on Facebook—and overwhelmingly negative.

Most negative comments came from users who seemed to be custodial parents, though a fair number also came from people who either appeared to be noncustodial parents or to identify with their perspective. As one might expect, their complaints almost always came from opposite points of view, with custodial parents typically voicing frustrations over a perceived lack of follow-through on enforcement in their particular case and noncustodial parents expressing a broader level of dissatisfaction with the program.

⁵¹ See section 2.III for a description of different types of engagement recognized in this cycle.

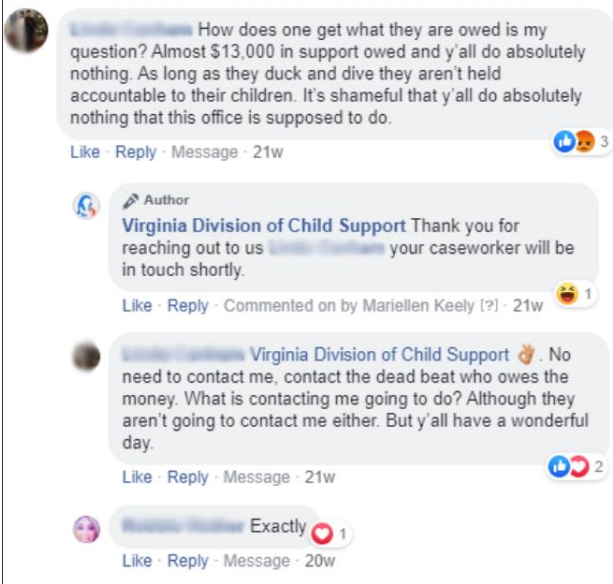


Figure 7: Top: Instagram post with negative comment. Bottom: Facebook post with a negative comment and an exchange with a DCSE staff member.

Since negative comments appeared almost as soon as the campaigns started, DCSE had to move quickly to address them, even before a formal policy had been finalized. Early on there had been some question about whether every negative comment needed to be addressed by phone with a DCSE representative. Eventually, DCSE adopted this as its formal policy. A sample of negative comment on Instagram post is included in Figure 7 above, along with one from a Facebook post that includes an exchange with a project staff member telling the person that their caseworker will contact them.

When parents sent that information and eventually talked with someone from DCSE, the resulting conversations succeeded in giving parents a chance to be heard. Yet it rarely led to

the kind of meaningful resolution that we had hoped to achieve, if only because the facts of most cases made it essentially impossible for DCSE staff to solve the underlying issue.

Whether successfully contacted by DCSE or not, most people only left one negative comment. The most persistent negative commenter, on the other hand, was eventually persuaded to talk to a member of the project staff and even appeared to be somewhat mollified over the course of the conversation—only to then return to posting one negative comment after another.

All that said, the negative comments received over the course of Cycle 2—amounting to at most a few dozen—should be seen in the context of the more than 500,000 people reached during that same time.

VI. Target Populations

As noted above in section 2.IV.E., all three intervention cycles target parents in Virginia who are eligible for DCSE services but not receiving them, with an emphasis on custodial parents who have also never participated in public benefit programs.

Because we decided to narrow our content focus in Cycle 2 to applications, the target audience for the second intervention does not include noncustodial parents, as we did in Cycle 1. Instead, this cycle considered its audience to be custodial parents or guardians in Virginia, including both those who have never received IV-D services from DCSE and those who have.

The location targeting we carried over from Cycle 1, however, winnow the audience to some degree. As explained above, in section 2.V.C.IV., we did not always succeed in restricting our audience to the location targets listed below in Table 3.

Table 3: Location Targets and Extended Focus Areas

Initial Target Area	Surrounding Areas Included in Ad Target Markets
Bristol (city)	Washington County
Brunswick County	
Caroline County	
Clarke County	
Colonial Heights (city)	Chesterfield County, Prince George County
Danville (city)	Pittsylvania County
Dinwiddie County	
Essex County	
Hampton (city)	Northampton County, Poquoson (city), York County
Henrico County	
Martinsville (city)	Henry County
Nelson County	
Newport News (city)	Isle of Wight County, James City County, Surry County, York County

Initial Target Area	Surrounding Areas Included in Ad Target Markets
Norfolk (city)	Virginia Beach (city)
Petersburg (city)	Chesterfield County, Prince George County
Portsmouth (city)	Chesapeake (city)
Richmond (city)	Chesterfield County
Roanoke (city)	Roanoke County
Salem (city)	Roanoke County
Staunton (city)	Augusta County
Suffolk (city)	Chesapeake (city), Southampton County
Waynesboro (city)	Augusta County
Winchester (city)	Frederick County

VII. Timelines

A. Overall Project Timeline

A broad outline of our project’s overall timeline appears below as Figure 8.




Figure 8: Overall project timeline.

B. Detailed Timeline for Cycle 3

The detailed timeline for Cycle 2 appears below as Table 4.

Table 4: Detailed Timeline for Cycle 2

Step	Plan Element	Approx. Date of Completion
1	Select social media platforms to be used for organic and paid social media campaigns	9/1/19
2	Develop budgeting tool to predict and track spending	10/1/19
3	Register social media accounts for organic and paid social media campaigns	10/23/19

Step	Plan Element	Approx. Date of Completion
4	Launch new Apply webpage	10/25/19
5	Ensure new online child support application ready to accept traffic from ads	11/1/19
6	Review 30 days of data and revise paid and organic ad strategies	12/2/19
7	Launch Lead Form ads	12/28/19
8	Review 60 days of data and revise paid and organic ad strategies	1/2/20
9	Implement corrected ad targeting	1/10/20
10	End Cycle 2 advertising	1/29/20

VI. Cycle 2 Outcome Measures

For Cycle 2 we intended to look at two types of outcome measures: one focused on online analytics and the other on caseload data. A list of outcome measures and sources, as well as the relevant dates, appears as Table 5 below.

Table 5: Outcome Measures for Cycle 1

Data Element	Frequency	Source	Applicable Dates
Click-through-rates (clicks divided by impressions) for Facebook ads	Cycle average	Facebook Ad Reporting	11/1/19-1/29/20 ⁵²
Click-through-rates (clicks divided by impressions) for Facebook ads	Cycle average	Facebook Ad Reporting	11/1/19-1/29/20
Click-through-rates (clicks divided by impressions) for Twitter ads	Cycle average	Twitter Ad Reporting	11/1/19-12/12/19 and 1/10/20-1/29/20
New cases opened by DCSE for child support services	Monthly average over same period of the previous calendar year	DCSE caseload data	11/1/19-1/29/20 versus 11/1/18-1/29/19

3. Results

I. Sample Size

We arrived at an estimate of the total number of individuals reached by Cycle 1 interventions by looking at the number of unique individuals reached by ads on Facebook. Known as *reach* in the online advertising world, this measure is only an estimate, since it attempts to account for duplicated or pseudonymous accounts. Still, we felt that this was the most reliable number for several reasons.

⁵² For reasons described in section II.V.C., there were multiple dates during Cycle 2 when Facebook and Instagram delivered almost no ads to our audience. We have nonetheless included these dates in our analysis.

- The paid campaign reached more people than the organic campaign, and since we cannot distinguish between the individual users reached by one or the other, we necessarily chose the larger one.
- We have no way of knowing if an individual user of one paid platform saw our ads on another. To avoid duplicating audiences as much as possible, we restricted ourselves to one platform.
- Twitter does not make data on reach readily available through its ad reporting tool. It also has a smaller user base.
- Facebook actively works to weed out accounts that do not appear to have been created by real people or organizations, as we experienced when trying to establish an account.
- While a branch of Facebook’s parent company, Instagram does not take the same stance toward fake or pseudonymous accounts. It also has a smaller user base.

Based on Facebook data, we reached approximately 536,000 unique users over the course of Cycle 2.⁵³

II. Results

We began this cycle with two research questions in mind:

1. Can social media marketing interventions drive positive public engagement with child support services?
2. Can promoting an online application portal through social media marketing increase the number of new applications for child support services?

Results we used to help us answer the first question appear in Tables 5 and 6. The data we used to address the second question appears in Table 7.

We used year-over-year data to gauge changes in the number of applications because such data is readily available and seemingly uniform in its definitions during the period studied. Such data does not, by itself, offer clear evidence that the variables we introduced during November 2019 through January 2020—that is, the online child support application and the social media advertising undertaken to support it—actually caused any of the changes in application numbers that we identify. Indeed, a number of potential variables could be at work here, including ongoing demographic shifts, changes in employment patterns, or even longer term changes in the reputation of the program that do not result from the activities of this project.

⁵³ Facebook explains how it defines reach here: <https://www.facebook.com/business/help/710746785663278>.

Still, these provide a broadly appropriate evaluative method of the sort that many program leaders might use in their day-to-day work.

Since DCSE had never placed social media ads in the past, no comparisons to previous results were possible. Instead, we used updated versions of the goals given in our revised Communications and Evaluation Plans, which rely on benchmarks generated by the company AdStage.⁵⁴ The basis of their estimates is proprietary, however, and in certain cases this creates ambiguities about what their measurement does and does not include. This ambiguity is particularly apparent when looking at click-through rates—our primary measure of physical digital engagement with the paid social media campaign in Cycle 2.

In Cycle 1 and on some social media platforms, clicks represent a fairly straightforward calculation of the number of times a user clicks on the destination link of an ad divided by the number of impressions that ad received.

Assessing click-through rates in Cycle 2 is somewhat more complicated, because social media platforms measure user engagement in ways other than simply clicking on a link. Under this engagement model, a click on an image or a click to an advertiser's profile page may also count toward a broad definition of click-through rate.

While we can clearly distinguish between these two measures, it is less clear which of these two options is being used in AdStage's benchmarks. Their simple reference to a click-through rate could be read as a sign that they mean the stricter version used in Cycle 1. Facebook and Instagram, on the other hand, emphasize the broader definition of click-through rate and their approach is therefore widely used in online advertising. Whether AdStage has adopted the same definition for those platforms is unclear, though, since they never define the term.

We will therefore give both in the sections for Facebook and Instagram, referring to them as strict (link clicks divided by impressions) and broad click-through rates (any type of click divided by impressions). Twitter does not classify multiple clicks in the same way as Facebook and Instagram; we therefore presumed our benchmark for Twitter describes only a strict click-through rate.

⁵⁴ Note that our goals for social media ads have been revised from our evaluation plan submitted in December 2019, which relied on third quarter 2019 data. In the interim our source for social media benchmarks, AdStage, released a set covering the fourth quarter of 2019—a period with a clearer seasonal overlap with Cycle 2. We have therefore updated our goals accordingly, while still following the basic approach of trying to achieve 11 percent above the industry average. For details see AdStage "Paid Media Q4 2019 Benchmark Report" (<https://cdn2.hubspot.net/hubfs/4350015/Benchmark%20Report/Q4%202019%20Paid%20Media%20Benchmark%20Report.pdf>).

Table 5: Outcome Measures and Results: Social Media Advertising Performance by Click-through Rate (CTR), Strict

Platform	Benchmark	Goal	Cycle 2 Strict CTR	Difference between Cycle 2 Outcome and Goal	% Difference between Cycle 2 Outcome and Goal
Facebook	1.14%	1.27%	0.94%	-0.33	-26%
Instagram	0.44%	0.49%	0.37%	-0.12	-25%
Twitter	2.36%	2.62%	5.44%	2.82	52%

Table 6: Outcome Measures and Results: Social Media Advertising Performance by Click-through Rate (CTR), Broad

Platform	Benchmark	Goal	Cycle 2 Broad CTR	Difference between Cycle 2 Outcome and Goal	% Difference between Cycle 2 Outcome and Goal
Facebook	1.14%	1.27%	1.74%	0.47	37%
Instagram	0.44%	0.49%	0.46%	-0.03	-6%

Table 7: Outcome Measures and Results: Search Ads (November 1, 2019–January 31, 2020 Compared to November 1, 2018– January 31, 2019)⁵⁵

Number of New Applications Opened for Child Support Services ⁵⁶	November 1, 2019– January 31, 2020	November 1, 2018– January 31, 2019	% Change
Total	8,979	9,368	-4%
Average per month	2,993	3,123	-4%

III. Analysis

A. *The Organic Social Media Campaign*

The paid social media campaign posted as often as we planned across all three platforms. Since all three accounts were new, they started with no followers. By the end of Cycle 2, DCSE’s Facebook account had approximately 200 followers, its Twitter account had approximately 90, and its Instagram account had more than 500. Such rapid growth seems remarkable, especially since we saw no evidence that the paid campaign contributed to the increase in followers on any of the accounts.

⁵⁵ DCSE does not report data on new cases by day, so we have used monthly data, which overlaps with the intervention cycle for all but two days.

⁵⁶ This measure revises an earlier one that appeared in our both our initial and revised evaluation plans, which focused on Never Assistance clients. Further analysis of the data led us to realize that we could not clearly separate out those cases in a way that was consistent with the definition used in OCSE’s 157 Report. We have therefore chosen to focus on all new cases, since any growth there is likely to come primarily from Never or Former Assistance clients, given the steady decline in Current cases.

B. Comments and the Hashtag Campaign

The hashtag campaign and the problem of how to respond to negative comments straddle both the organic and paid campaigns.

Virtually every post we made across both the organic and paid campaigns used the hashtag we created, so the campaign in that limited sense was successful. Yet the goal of using the hashtag to galvanize a wider movement that could celebrate the program or its goals never materialized. Indeed, searches conducted on Facebook, Instagram, and Twitter during and after Cycle 2 returned no results for #supportVAkids that did not either originate with this project or come from someone associated with it, whether as part of the project team or as a peer from another jurisdiction brought in for guidance.

As with the comments, the apparent lack of enthusiasm for a hashtag offering a positive take on child support probably reflects wider views of the program. Perhaps future interventions could improve on these results by introducing more varied organic content and connecting it to the hashtag campaign, rather than solely restricting it to child support. We could have also boosted both the hashtag campaign and our overall levels of demonstrated engagement by connecting our posts more often with existing hashtags, particularly on Instagram, where both organic and paid campaigns routinely combine a dozen or more hashtags. As it is, we used only our #supportVAkids hashtag.

C. The Paid Social Media Campaign

The outcomes of the paid campaigns in Cycle 2 suggest that DCSE and other child support programs may not see clear benefits to broad-based social media advertising. Nonetheless, certain results, examined in a wider context, do suggest benefits to this method of digital marketing.

In contrast with Cycle 1, where we consistently achieved above our benchmarks and goals, we fell below them in Cycle 2 for both Facebook and Instagram, where our strict click-through rates were 26 percent and 25 percent, respectively. We came closer to our goals if we presume that the underlying benchmark represents a broad click-through rate for those platforms, with Facebook 37 percent over the goal and Instagram just 6 percent below it. (Measuring this way, Instagram was also slightly above its industry-wide benchmark.)

Performing even close to an industry-wide benchmark, one overwhelmingly oriented toward the private sector, should not to be taken for granted, however, especially since we could not build on the usual foundation of a successful paid social media campaign: a well-established organic social media presence. Given these mitigating factors, we tend to view the results of our paid campaign as a sign that much of what we were doing was working—just not enough to

surpass the level established by stores, spas, and other entities able to offer a simpler and less fraught statement of value than a child support program reasonably can.

Another way of looking at the performance of the paid campaign, though not formally included as a goal for this cycle, is cost per conversion. One limitation for the data in Cycle 2 is that the ad platforms we used failed to record any direct conversions from individual ads, other than those associated with the Lead Form ads described in section 2.V.C.IV.⁵⁷ Nonetheless, some broad points can still be made about how conversion functioned in Cycle 2 social media ads relative to the Google Search and Display ads used in Cycle 1.

Before comparing those cycles, we should acknowledge the differences between them. In the last cycle we had three sets of Search ads and three sets of Display ads all pointing to a single page (the main project webpage) and with a single conversion goal: completing a short contact form to request a specific type of information. All nine ad sets were delivered by a single provider (Google) but, on the Display side, across hundreds of distinct sites and platforms.

In Cycle 2, we had two ad providers (Facebook and Twitter) delivering three interrelated strains of advertisements (branding and awareness, conversions, and Lead Form) that pointed users to distinct conversion points (the main project webpage, the Apply page, and a Lead Form window within either Facebook or Instagram). Though all three ad sets shared the same overall goal (encouraging users to apply for child support), the visitor to the main project webpage had to complete the same contact form from Cycle 1 to be counted as converting, while users on the Apply page converted by clicking on a button marked “Apply Now!” that led them directly to the application.

Conversion on the contact form thus required a higher level of effort than with the Apply page or Lead Form ads, since users had to both type in their contact information (at least some of which would have been prepopulated on the Lead Form) and then hit a button to send it (all that was required on the Apply page). The cognitive or emotional bar for converting after viewing one of the ads from Cycle 2, however, would seem to be higher, at least for the conversion ads that sent users to the Apply page, since each user would have to feel confident in their knowledge of what it means to apply and have the motivation to at least access the application directly, if not necessarily submit it. Our sense that this part of Cycle 2 conversions required higher levels of cognitive or emotional readiness than in Cycle 1 is purely suppositional; we did not undertake the surveys or biometric measurements necessary to

⁵⁷ We have not been able to determine why the analytics engines available through Twitter and Facebook could not show this data. Regardless, we hope to correct it in Cycle 3 or at least be able to offer a set of platform-specific conversion-related data (perhaps using Google Tag Manager) for the final evaluation.

assert this with greater confidence than that. But it seems worth noting it as a potential influence on the outcomes.

Current data on conversions in Cycle 2 appears below in Table 8. Overall, we recorded 355 conversions in Cycle 2 across the three venues (the main project website, the Apply page, and the Facebook Lead Form). This averages out to just under 4 per day across the 90-day cycle at an average cost of \$64.75 per conversion. While higher than the \$13.44 paid per conversion for Google Search ads in Cycle 1, the overall average cost per conversion in Cycle 2 remains far lower than the \$546.52 per conversion for Google Display ads in Cycle 1.

Table 8: Cost per Conversions by Platform, Cycle 2

	Cycle 2 Conversions across All Platforms	Ad Spend by Conversion Target			Total Cost per Conversion
		Twitter	Facebook	Total	
Apply Page (apply.supportvakids.com)	182	\$956.86	\$7,893.87	\$8,850.73	\$48.63
Main Target Page (supportvakids.com)	136	\$3,523.96	\$8,895.22	\$12,419.18	\$91.32
Total of Apply and Main	318	\$4,480.81	\$16,789.09	\$21,269.90	\$66.89
Lead Form	37	NA	\$1,714.93	\$1,714.93	\$46.35
Grand Total	355	\$4,480.81	\$18,504.02	\$22,984.83	\$64.75

Calculated as percentage of the number of conversions (355) divided by the total number of link clicks across all platforms (40,594), the conversion rate for Cycle 2 was 0.87 percent. This is lower than the 9.21 percent conversion rate across all industries and the 5.60 percent rate for legal services providers given by Wordstream, the only widely cited source we could find with a benchmark for this metric for Facebook ads.⁵⁸

A majority (51 percent) of those conversions were to the Apply page. If one accepts that the costs of those conversions are equivalent to the cost of ads directing users specifically to that page, then the average cost of conversion for those ads was \$48.63. This logic is somewhat at odds with the overall focus in Cycle 2 of sending users down an advertising funnel that begins with branding and awareness ads (which had the main project page as their conversion point) and ends with this point of conversion to the Apply page. Yet even if one makes the Apply page ads responsible for the cost all of the ads in Cycle 2, the cost per conversion goes to \$126.29, a figure that is still less than a fourth of the average cost per conversion for our Google Display ads in Cycle 1. The end point of this conversion also positions users to apply for child support in

⁵⁸ Mark Irvine (2020). "Facebook Ad Benchmarks for YOUR Industry." Wordstream. <https://www.wordstream.com/blog/ws/2017/02/28/facebook-advertising-benchmarks>.

a matter of minutes, not the minimum of a few days that is required when asking for information on the main project page contact form.

About 38 percent of the remaining conversions were on the main project webpage. The average cost per conversion for those ads was \$91.32—somewhat less than twice the average cost per conversion for the Apply page. Since ads linking to the main project webpage received two-thirds (66 percent) of our ad impressions on Facebook, Instagram, and Twitter,⁵⁹ the lower number of conversions associated with the main target webpage is not simply a function of the number of times the ads were shown. Instead, we propose three possible reasons:

- Ads for the Apply page tended to (but did not always) feature longer videos, since they were farther down the tunnel and we felt they would be more receptive to more in-depth messaging. Perhaps longer videos are more persuasive.
- The call-to-action for ads linking to the main page usually read “Learn More,” while Apply page ads tended generally used “Apply Now.” Perhaps the more active and narrowly directive call-to-action is more likely to lead users to a specific outcome.
- The ads to the Apply page were the end of the advertising funnel, so users who saw the earlier ads should have been persuaded by those ads and open to the guidance given in the later ones.

While we have no firm evidence to support any of these possibilities—and acknowledge that other factors could also be driving this outcome—we at least see the last of these possibilities as potentially affirming the value of planning ads using a funnel model.

One additional factor that might have affected the performance of our ads was the timing of the cycle itself. Starting on November 1, just as the ads for the winter holidays begin, and continuing through to the end of January 2020, shortly before the National Football League’s Super Bowl game, Cycle 2 fell during a time of heavy advertising, particularly for consumer goods.

One effect of this might have been that our ads seemed out of place with the season. Amid ads for toys, clothing, and electronics sprinkled with nods to winter and specific religious traditions, parents might have seen our ads as stressful reminders of a lack of income for their children that could bring with it divisiveness at odds with themes of family and togetherness characteristic of Thanksgiving, Hanukah, Christmas, New Year’s, and other holidays during this time.

⁵⁹ This 66 percent figure holds true even if impressions from the Lead Form ads are included in the overall calculation.

Whether these subjective considerations tend to make parents less inclined to apply for child support or whether larger structural factors (such as end of the year bonuses or anticipated tax refunds) are at play, DCSE has seen some seasonal effects in the opening of child support cases over the last few years. As demonstrated below, in Table 9, the number of new cases opened in January, November, December of the calendar years 2016 through 2019 fell below or roughly at the same level as the median for that year. The deviation from the median becomes clearer when reading backwards from January to the previous calendar year median. For example, the 3,611 cases opened in January 2017 falls quite short of the median of 4,456 for 2016, and so forth. The differences between the last two months of the calendar year and the median seem to have sped up in the last two years, as seen in the figures from November and December 2018 and 2019. A complete illustration of DCSE’s newly opened cases by month appears below, as Figure 9.

Table 9: Total New Child Support Cases Opened for DCSE during January, November, and December, 2016–2019

	January	November	December	Calendar Year Median
2016	3,685	4,402	4,505	4,456
2017	3,611	4,143	3,745	4,135
2018	3,337	3,144	2,876	3,675
2019	3,348	3,066	2,763	3,330




Figure 9: Total number of newly opened DCSE child support cases, by month, between 2016 and 2019.

Another element apparently affecting the results of our paid campaign, if only in terms of cost, was an apparent increase in prices for ads on social media, as seen in changes to the average amount we paid for an individual user to click one of our ads (known as cost per click). Cost per click is admittedly not a pure indicator of the overall ad volume relative to the number of users seeing the ad, but it does provide a method for evaluating the amount of competition advertisers face in a given period or for a given type of content.⁶⁰

Figure 10, below, shows total daily spend on social media across all three paid platforms (Facebook, Instagram, and Twitter) and the average daily cost of a click—the user action that established the price for most of our social media ads—over the course of the 90-day intervention.⁶¹ After a sharp drop a few days into the cycle, the cost per click moves upward over the course of November 2019, hitting peaks near the Thanksgiving holiday, before descending in January 2020—a series of shifts that is largely independent of spending in November and most of December. The relationship between spending and cost per click in January is complicated by our return to advertising on Twitter on January 10 and the immediate change in performance that increased clicks dramatically and led to an equally intense drop in the cost per click.

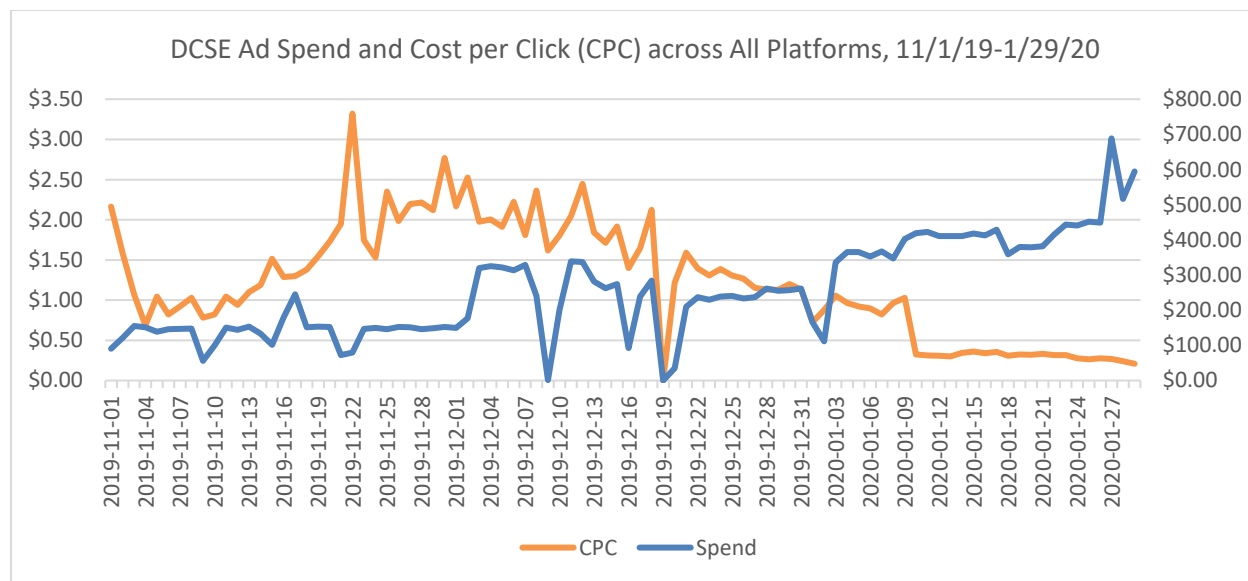


Figure 10: Total ad spending across all DCSE social media platforms (Facebook, Instagram, and Twitter) in Cycle 2 compared to cost per click (CPC). The sharp decline in cost per click beginning on January 10, 2020 is a result of the return to advertising on Twitter and the lower cost per click associated with that segment of the Twitter campaign.

⁶⁰ Our interim evaluation for Cycle 1 noted, for example, the relatively high cost per click for Google Search ad for legal services—a consequence of high demand for a limited amount of space available in search results.

⁶¹ This figure averages the daily cost per click by totaling all the clicks received on the platforms and then dividing them by the total cost of ads for that day. Another method would be to average the daily cost per click across all three platforms but that would provide more variable day-to-day results, particularly during the first month, when Facebook’s algorithm was still determining ad placements.

By comparison, Figure 11 isolates the cost per click and cost per thousand impressions on Facebook and Instagram. Since the performance of our ads on these platforms never went through the kind of sharp changes we experience on Twitter, they perhaps give a clearer representation of possible seasonal changes affecting our performance. As the figure shows, the same overall upward movement in the cost per click from early November to mid-December is followed by a gentler decline and then a similar upward shift as January comes near a close. Meanwhile, the cost per thousand impressions—a cost that largely reflects the cost per click in the bidding model we used—follows a similar contour but with fewer marked fluctuations.




Figure 11: Cost per thousand impressions (CPM) and cost per click (CPC) across DCSE's paid Facebook and Instagram campaigns in Cycle 2.

Monthly average performance for ads on Facebook and Instagram in each of those months suggests a similar pattern. Our average daily cost per click on those two platforms was \$1.08 in November, \$1.45 in December, and \$1.35, in January, and our average cost per thousand impressions was \$9.31, \$12.66, and \$8.74, respectively.

These increases in costs per click toward the end of the calendar year seem in line with data on industry trends in 2018, at least for Facebook. According to the website AdEspresso the average cost per click for the first two quarters of 2018 was \$0.09. It then rose to \$0.12 for the third quarter before jumping to \$0.20 for the fourth.⁶² A graph of the AdEspresso data on the average cost per click per month on Facebook in 2018 appears below as Figure 12.

⁶² Tory Wenger (2019). AdEspresso "Understanding Facebook Ads Cost – 2018 (& 2019) Benchmarks!" <https://adespresso.com/blog/facebook-ads-cost/>.

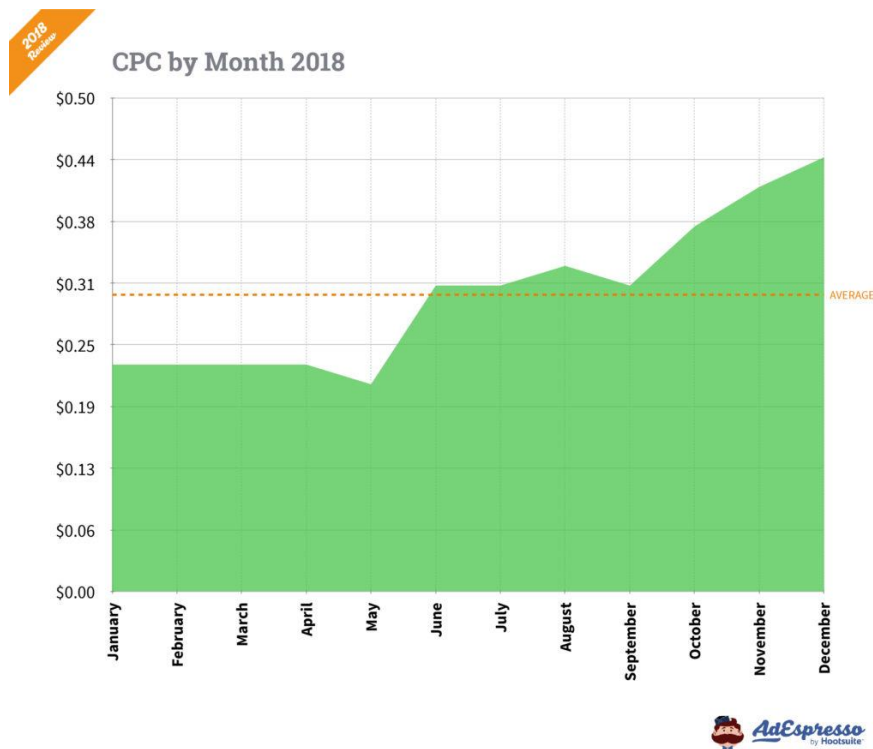


Figure 12: Average cost per click (CPC) by month on Facebook in 2018. Source: Tory Wenger (2019). AdEspresso “Understanding Facebook Ads Cost – 2018 (& 2019) Benchmarks!” <https://adespresso.com/blog/facebook-ads-cost/>.

Our results on Twitter show similar seasonal effects in November and December 2019 but not in January. One factor contributing to the differences between Twitter and the other platforms may be related to the overlap of Cycle 2 with a period of heavy advertising on the part of presidential campaigns, primarily but not exclusively on the part of Democratic candidates.⁶³ While Facebook and Instagram allowed political advertisements throughout Cycle 2—and will apparently continue to do so for the foreseeable future—Twitter announced two days before the start of Cycle 2 that it would extend its earlier ban on issue ads in the US to a global ban on political ads.⁶⁴ The policies did not take effect, however, until November 15, 2019, a little more than two weeks into our second cycle.⁶⁵

⁶³ The campaign for then-Democratic primary contestant New York City mayor Michael Bloomberg spent \$8.2 million on digital marketing alone in the fourth quarter of 2019. Source: Sarah Almukhtar, Thomas Kaplan and Rachel Shorey (Feb. 1, 2020). *New York Times*. “2020 Democrats Went on a Spending Spree in the Final Months of 2019.” <https://www.nytimes.com/interactive/2020/02/01/us/elections/democratic-q4-fundraising.html>.

⁶⁴ Twitter’s 2018 ban on issue ads in the U.S. is available here: Del Harvey and Bruce Falck (Aug. 30, 2018). Twitter. “Announcing New US Issue Ads Policy.” https://blog.twitter.com/en_us/topics/company/2018/Announcing-new-US-issue-ads-policy.html. The October 30, 2019, global ban on political ads was announced in a threaded tweet from Twitter’s Chief Executive Officer, Jack Dorsey: <https://twitter.com/jack/status/1189634360472829952>.

⁶⁵ Twitter’s current statement of its policy for political ads is available at <https://business.twitter.com/en/help/ads-policies/prohibited-content-policies/political-content.html>. For an early analysis of the differences between how Facebook and Twitter treated political ads, see Emily Stewart (Nov. 15, 2019). *Vox*. “Twitter Is Walking into a

Seasonal advertising effects on Twitter are also obscured by the significant shift between what our first Twitter campaign (lasting from November 1 to December 12, 2019) and our second (lasting from January 10 to January 29, 2020, after we corrected the demographic and geographic targeting of the ads).⁶⁶




Figure 13: Cost per click (CPC) and strict click-through rate (CTR) for Twitter paid campaign in Cycle 2.

When results from both Twitter campaigns are combined, ads on that platform achieve a click-through rate of 5.44 percent and an average cost per click of \$4.56. But these averages are skewed by the stark contrast in performance between the two campaigns. While the first campaign had an average click-through rate of 0.15 percent and an average cost per click of \$6.74, the average click-through rate for the second Twitter campaign was 11.93 percent, with an average cost per click of just \$0.08. The contrast is illustrated above in Figure 13, which also shows the overall rise in the average cost per click over the course of November and into the middle of December.

However, as noted above, in section II.V.C.IV, we came to believe that the results of this second Twitter campaign were consistent with activities initiated by a click bot rather than real users.

Minefield with Its Political Ads Ban.” <https://www.vox.com/recode/2019/11/15/20966908/twitter-political-ad-ban-policies-issue-ads-jack-dorsey>.

⁶⁶ A small number of ads were delivered by Twitter on December 13. Because they represent only a fraction of a day’s results and performed out of line with the rest of our campaign, we have not included them in an analysis of the first segment of our Twitter campaign.

While the only evidence we can present for this is circumstantial, our analysis aligns with other published descriptions of click fraud on Twitter.⁶⁷

The three most important pieces of data that support our argument for click fraud are:

1. The nearly one-to-one relationship between link clicks and engagement.
 - During the first part of the campaign (November 1 to December 12, 2019) we received only 347 clicks on the link to our website (termed “link clicks”) but 4,304 types of physical engagement overall (including link clicks, clicking on images, clicking through to view our profile, and so forth). This demonstrated a more than 12:1 ration of engagements to link clicks.
 - In the second part (January 10 to January 29, 2020) we received 25,623 clicks and 25,773 engagements, for a nearly 1:1 ratio. Such a dramatic shift seems unlikely in a short time.
2. A much larger percentage of clicks for the second campaign came from accounts that include languages other than English.
 - In the first campaign, Spanish speakers made up 1.44 percent of clicks, followed by Chinese speakers at 0.96 percent. This seems in line with our expectations and with Census data on languages spoken at home in Virginia, which lists languages in exactly that same order of prevalence: English, Spanish, and Chinese.⁶⁸
 - In the second campaign, Spanish speakers were 21.48 percent of our clicks and our fourth-largest audience, after English, Japanese, and Arabic speakers. Chinese speakers, meanwhile, gave us almost the same number of clicks that they did in the last campaign.
 - All of these numbers for the second campaign are outside of what we would expect based on Census data percentages for Virginia.
 - Overall, in the first campaign speakers of languages other than English made up only 4.03 percent of the total number of clicks. In the second, they made up 113.42 percent.⁶⁹

⁶⁷ See, for example, a blog post from the advertising firm Penna Powers titled “Twitter’s Bot Problem is Finally Out of the Bag”: <https://pennapowers.com/twitters-bot-problem-is-finally-out-of-the-bag/>.

⁶⁸ See Statistical Atlas (2019). “Languages in Virginia.” <https://statisticalatlas.com/state/Virginia/Languages>.

⁶⁹ Numbers total more than 100% because accounts can have multiple languages associated with them so they get counted multiple times.

3. The amount we spent on ads in the second campaign skewed toward accounts that include languages other than English.
 - In the first campaign, 4.60 percent of our total spend went to accounts that included a language other than English.
 - In the second campaign, slightly more than 100 percent of our spend went to those kinds of accounts.

The final and most important support for our view that these clicks aligned with a pattern of click fraud came from the realization that virtually all of the link clicks in our second Twitter campaign came from its third-party ad distribution network, the Twitter Audience Platform. Since these third-party sites receive a percentage of each click, they could be seen as having a financial incentive to drive up the numbers artificially.

D. Online Child Support Application

We also puzzled over the interpretation of some of the data related to the online child support application. Even before we started sending users to the Apply page and from there to the online application, users were submitting child support applications through the new portal. Not long after we started, the number of online applications increased dramatically—but far out of proportion to the number of users recorded as converting from our ads. Indeed, by the end of Cycle 2 we had recorded only 182 conversions from the Apply page while approximately 1,900 applications had been submitted during that same period.

The cause for this disjuncture is unclear. It might come from the application’s prominent placement on the MyChildSupport page, which is often the first search result for common searches such as “Virginia child support.” It might also be the result of court or program staff being made aware of the online application and sending parents to that rather than encouraging them to complete the paper application. A third possibility is that the branding and conversion ads placed over the course of Cycle 2 succeeded in planting the idea of applying, even if no specific ad caused them to convert.

We may also be asking too much of any marketing campaign to, by itself, reverse a phenomenon as complex as the decline in applications for child support services. Our overall sense is that individuals are driven to apply for child support because of significant life events: a sudden change in a child’s needs, the end of another parent’s informal support, or the threat of homelessness, among many other causes. The relatively gentle persuasiveness of a short video might bring someone to buy a new toy for their child or a pair of shoes for themselves, but perhaps the decision to seek out child support is too significant to be meaningfully affected by something so slight.

Regardless, we did see tentative evidence that the combination of advertising and the online application may have helped reduce the speed with which Virginia—like other IV-D programs—is seeing the number of applications decline. Looking at the three-month period between the beginning of November 2017 and the end of January 2018, Virginia’s IV-D program saw the average number of new cases opened each month decline by 431 (10 percent) over the same period in 2016 and 2017.⁷⁰ That decline continued during the same period between the last two months of 2018 and the first month of 2019 year, with a reduction of an average of 619 cases per month (17 percent). Between November 2019 and the end of January 2020, however, the average number of cases opened per month declined by only 130 (4 percent)—a notable deviation from the trend of the previous two years, albeit not one inarguably tied to the social media campaigns of Cycle 2 or the new online application.

Still, these changes in case trends suggest that digital engagement methods—including a more contemporary approach to requesting child support services—may at least help to reduce the program’s decline, even if it cannot reverse it entirely.

4. Lessons Learned and Next Steps

I. Lessons Learned

We see three primary takeaways from Cycle 2.

- 1) **Each social media platform brings both risks and rewards.** The mixed outcomes of this cycle suggest that no social media platform offers a clear path to strong, positive social engagement, regardless of whether one comes to them as an organic presence or a paid advertiser. At the same time, we saw positive outcomes from each platform: the broad reach with Facebook, the impressive (if not entirely trustworthy) click-through rate for ads with Twitter, and the increasing number of users following DCSE on Instagram. Child support programs thus have reason to approach each of them with some wariness but perhaps a greater amount of optimism, if they are able to stick with them and keep testing platform-specific strategies to improve engagement.
- 2) **Tools matter for complex advertising projects spanning multiple platforms and dozens of paid ads.** The problems we encountered with the incorrect targeting on our ads and with gaps in ad delivery on Facebook and Instagram came in many ways and from different sources. The incorrect targeting was partially human error. The gaps in ad delivery were caused to some degree by long delays in ad approvals from Facebook. One thing that perhaps could have reduced the likelihood of both, however, would have been an efficient tool that could have allowed us to easily transfer the ad information to all three platforms at once (rather than having to enter it manually) and to easily reuse

⁷⁰ We are treating new cases opened as a proxy for the number of applications submitted, since the latter is not normally tracked by DCSE. While we recognize that the two are not perfectly aligned, they seem closely correlated enough to allow for the kind of broad, tentative inferences made here.

and review the targeting information associated with each ad across all three platforms. While we know such tools exist, none seems priced at the right level for a child support program to easily take advantage of without significantly reducing the efficiency of its ad spending. We will continue to investigate these for Cycle 3.

- 3) **Slowing the decline in applications may be the best that programs can hope for when it comes to addressing the gradually diminishing caseloads seen in child support programs across the country.** While we had hoped to see the number of applications for child support increase as a result of this intervention, we have come to accept that all digital engagement may be able to do is slow the decline. If that's the case, that remains a good argument for taking up these strategies, since the costs remain relatively small compared to the costs of staffing or many technological solutions.

II. Next Steps

We plan the following set of interventions for Cycle 3:

- 1) **A test of a "best practices" advertising model.** We want to see whether a paid campaign centered around Google Search but supported with remarketing advertisements delivered by Google Display, Facebook, Instagram, and Twitter ads can offer a strong return on investment. The ads will drive users to the main project website and its contact form rather than the Apply page
- 2) **Snapchat.** We plan to create a special set of ads and post them on a weekly basis to Snapchat to see whether we can broaden awareness of the program among the younger demographic associated with that platform.
- 3) **Webinar for family law attorneys.** We will administer a webinar for family law attorneys in Virginia to encourage them to send their clients to us for child support. We hope to receive approval from the Virginia Bar Association to give Continuing Legal Education (CLE) credits to attendees.
- 4) **Podcast.** We plan to produce a series of podcasts on child support responding to various common misconceptions about the program.
- 5) **LinkedIn.** We will be using LinkedIn ads to promote our legal webinar to family law attorneys in Virginia.

The outcomes of Cycles 1 and 2 have informed a number of elements of their implementation, including:

- 1) Continuing to use the location targeting used throughout Cycle 1 and partially in Cycle 2.
- 2) Continuing to use behavioral concepts in crafting the language for Cycle 3 ads and for the podcasts and webinar.
- 3) Reusing text from Cycles 1 and 2 and videos from Cycle 2 for ads on Facebook, Instagram, and Twitter.

- 4) Continuing to use a dynamic budgeting model that starts ad purchases with relatively small budgets and increases them incrementally based on performance.
- 5) Continuing to analyze case openings to see whether we can discover a stronger causal connection between them and the digital engagement project as a whole (including, for these purposes, the electronic application).

Schedule of Appendixes

[A: Apply Webpage](#)

[B. Online Contact Form](#)

[C. 50 Most Productive Search Terms \(Ranked by Total Number of Clicks\)](#)

D. Map of Contact Form Responses [included as a separate PDF]

Appendix A: Apply Webpage

The final text of the Apply page (<https://www.dss.virginia.gov/supportvakids/apply/>) is given below, with a simplified markup (h1 = Heading level 1, h2= heading level 3, etc.) included to indicate relative importance. Other formatting is used to indicate rough placement of the text.

An image of the page appears above as Figure 3.

<h1>Applying for Child Support in Virginia Has Never Been Simpler</h1>

As Virginia's Division of Child Support Enforcement (DCSE), we want children to receive support. And we know that to do that, we have to make child support simpler for everyone.

That's why we built a new online application into our MyChildSupport portal. Now you can apply for child support on your phone, computer or tablet, 24 hours a day.

<h2>Your child's needs just can't wait</h2>

<button>APPLY NOW!</button>

No registration required. We respond within two business days.

<h3>More Services for Custodial Parents (Parents Wanting to Receive Child Support)</h3>

1. NEW 1-page child support application
2. Customer service experts trained to walk you through the process
3. Only \$35 annually (and only if we collect more than \$550/year on that child support case)

Don't lose out on the money your child needs!

Appendix B: Lead Form Ad

The final text of the Lead Form ad appears below. An image from a draft form is given above as Figure 6.

<h1>Calling all Moms and Dads!</h1>

- Child support orders always include medical support
- No fee for new applicants
- Child support is linked to kids doing better in school

Complete the form and we'll respond within two business days.

Full name [typically auto-populated]

Email [typically auto-populated]

Phone number [typically auto-populated]

What can we do for you?

- Open a child support case with VA Social Services
- Get help with an existing child support case
- Learn about our employment, co-parenting, and mediation services
- Connect to low-cost paternity testing
- [Other]

<button>Done</button>

Appendix C: Ad Performance in Cycle 1 versus Cycle 2 (Partial), Ranked by Spends in Top-Ten Designated Market Areas (DMAs)⁷¹

Table 1: Intervention Cycle 1: Google Search and Display Advertising
% of Total

Market Area (User Location)	Ad Spend	Cost	CTR
Norfolk-Portsmouth-Newport News VA	\$8,043	41%	0.62%
Richmond-Petersburg VA	\$5,276	27%	0.66%
Roanoke-Lynchburg VA	\$2,622	14%	0.64%
Washington DC (Hagerstown MD)	\$1,191	6%	0.53%
Harrisonburg VA	\$656	3%	0.62%
Tri-Cities TN-VA	\$625	3%	0.52%
Charlottesville VA	\$217	1%	0.48%
Raleigh-Durham (Fayetteville) NC ⁷²	\$141	1%	0.55%
Greensboro-High Point-Winston Salem NC	\$129	1%	0.48%
New York, NY	\$53	0%	0.92%

Table 2: Intervention Cycle 2 (November 1, 2019–January 4, 2020): Facebook and Instagram
% of Total

Market Area (User Location)	Ad Spend	Cost	CTR
Norfolk-Portsmouth-Newport News VA	\$2,684	27%	0.80%
Richmond-Petersburg VA	\$2,010	20%	0.75%
Roanoke-Lynchburg VA	\$1,831	18%	0.89%
Washington DC (Hagerstown MD)	\$1,589	16%	0.91%
Tri-Cities TN-VA	\$559	6%	0.87%
Non-DMA Region ⁷³	\$476	5%	0.41%
Harrisonburg VA	\$332	3%	0.96%
Charlottesville VA	\$135	1%	0.80%
Bluefield-Beckley-Oak Hill WV	\$102	1%	1.02%
Raleigh-Durham (Fayetteville) NC	\$99	1%	0.61%

⁷¹ DMAs are a proprietary grouping, established by the Nielsen Company, of 210 areas in the U.S. They are named by their largest urban areas but include suburbs, rural areas, and so forth.

⁷² The lists includes areas that are nominally outside our geographic targets in part because the DMAs overlap with our targets (as with the Raleigh-Durham area, which takes in part of southern Virginia) and in part because some ads were served in those areas, albeit only a small number, as the figures for New York City suggest.

⁷³ While “non-DMA Region” clearly indicates an area that falls outside the normal boundaries established by the Nielsen Company, it is not clear whether this encompasses other countries or simply parts of the U.S., such as rural Alaska, that other sources describe as non-DMA areas.

Table 3: Intervention Cycle 2 (November 1, 2019–January 4, 2020): Twitter

Market Area (User Location)	Ad Spend	Cost % of Total	CTR
Washington DC (Hagerstown MD)	\$883	38%	0.12%
Norfolk-Portsmouth-Newport News VA, US	\$555	24%	0.14%
Richmond-Petersburg VA	\$383	17%	0.15%
Roanoke-Lynchburg VA	\$360	16%	0.17%
Harrisonburg VA	\$75	3%	0.13%
Tri-Cities TN-VA	\$28	1%	0.12%
Charlottesville VA	\$13	1%	0.03%
Bluefield-Beckley-Oak Hill WV	\$0	0%	0.17%
New York NY	\$0	0%	0.00%
Raleigh-Durham NC	\$0	0%	0.00%

Appendix D. Contact Form Responses

See following pages for maps of contact form responses during Cycles 1 and 2 and over the course of the entire period from June 2019 through the end of January 2020.