

Cost Efficiency Analysis: Unconditional Cash Transfer Programs

The IRC's unconditional cash transfer programs range in cost efficiency from 14 cents for every dollar transferred up to \$1.32 for every dollar transferred, with programs that reached more households using community-based targeting achieving the highest cost efficiency.

With the unprecedented scale of emergencies facing the humanitarian community, we are in urgent need of programs that enable households to meet their basic needs and avoid negative coping mechanisms, and which use our scarce resources efficiently. Unconditional cash transfer programs—in which beneficiaries receive cash or vouchers that they can use as they please, rather than in-kind food or shelter assistance—are rapidly gaining popularity as a means to reach large numbers of people in need. Not only are cash transfers generally cheaper in administrative cost per dollar of value transferred¹, they aim to give beneficiaries greater dignity and control to prioritize their own needs. Nonetheless, there are still many questions being worked out about the optimal method to target and deliver cash transfers in different environments. One aspect to understand about these different program design options is their cost implications—will some program design choices dramatically increase or decrease the efficiency with which we transfer money?

The IRC is committed to maximizing the impact of each dollar spent to improve our clients' lives; our Best Use of Resources (BUR) Initiative specifically focuses on improving the reach and impact of IRC programs, by using cost analysis to identify the most cost efficient and cost effective ways to deliver programs. In this analysis, we have examined 8 unconditional cash transfer programs across 7 countries, using existing financial and administrative data to produce an estimate of each program's cost efficiency.

- **Cost efficiency analysis—estimating how much it costs in administration and program management per dollar transferred to beneficiaries with different program designs—helps decision-makers to make the best use of available resources.** The “cost transfer ratio” is the ratio of all non-transfer costs, such as staff time, targeting surveys, or transfer fees, to the value of the money that was transferred to recipients throughout the program. Comparing this ratio across programs allows us to see how design choices about the amount transferred, the scale of programs, and the targeting and transfer method affect a program's cost efficiency.
- **The IRC's unconditional cash transfer programs have a wide range of cost efficiency, from a minimum of 14 cents for every dollar transferred up to \$1.32 for every dollar transferred.** While some IRC programs are extremely cost efficient, there are others which cost more in administration than the money they are giving away, and we can learn from this analysis what design choices might make our work more cost efficient.
- **The biggest single factor driving cost efficiency is the scale at which programs are run—reaching more households spreads the fixed costs of country support over a wider pool of beneficiaries, driving down per-household costs dramatically.** The IRC may want to consider guidelines that cash transfer programs should reach some minimum number of households in the future.
- **Program design choices about targeting method have more of an impact on cost efficiency in contexts where there is a large difference between local and international wage levels.** In contexts where the price level is very low and fewer dollars get transferred to each beneficiary, non-transfer costs take up proportionally more of a program's total costs than in contexts with high price levels and larger transfers. In low-price contexts, the cost of giving money to a wider pool of beneficiaries and accepting some margin of error may actually be lower than the cost of extensive targeting activities.

¹ Doocy, Shannon, and Hannah Tappis. "The Effectiveness and Efficiency of Cash-based Approaches in Emergencies: A Systematic Review." *The Campbell Collaboration* (2015).

Cash Transfers at the IRC: What, Where, and How?

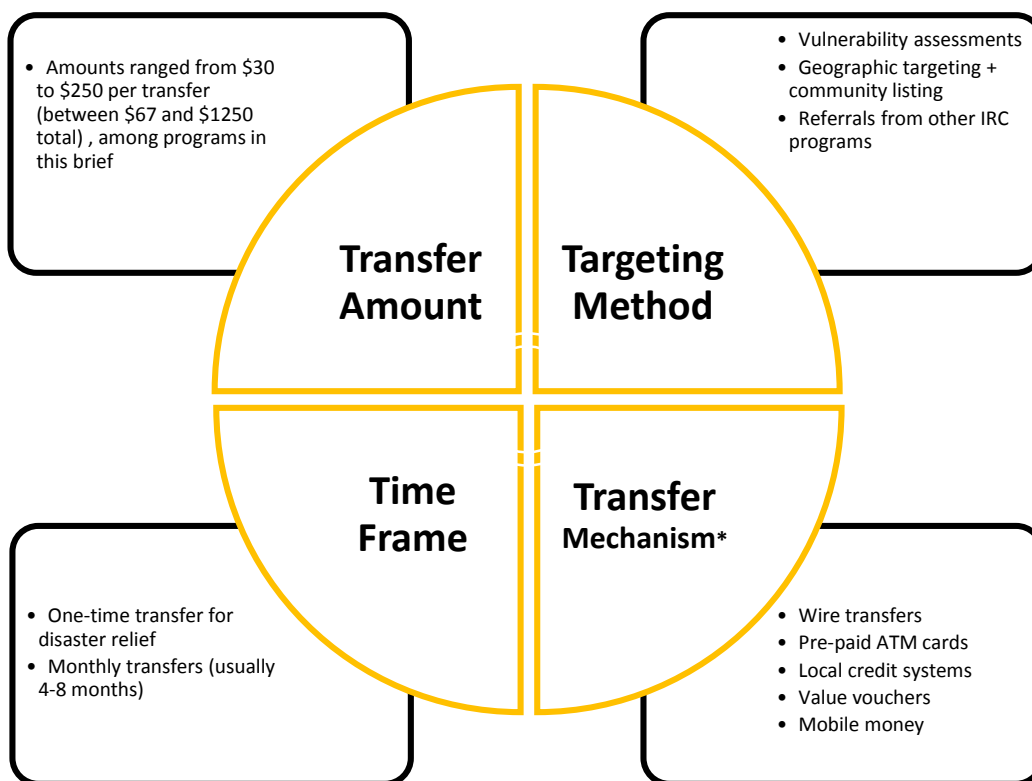
This analysis focuses on estimating the cost efficiency of the IRC’s unconditional cash transfer programs, and exploring which factors help such programs achieve maximum cost efficiency. The phrase “transfers” can encompass several different type of interventions which, although they are similar, have distinct advantages and drawbacks. The IRC’s wider cash transfer work falls into several categories:

- Cash for work (CfW) where transfers of different kinds are conditional on work,
- Commodity vouchers: Vouchers for specific items like food items or winterization kits,
- Value vouchers: Vouchers that can be used to purchase any items at participating stores, and
- Unconditional cash transfers, in the form of actual cash, wire transfers, or credit through local systems

The distinction between types of transfers are important, especially since different delivery systems and types of conditionality may be more or less cost efficient per dollar transferred to beneficiaries. The organization’s current emphasis on unconditional cash transfers is driven in part by the idea that such transfers are more cost efficient than conditional or in-kind transfers, because they require less complex supply chains and incur fewer transport costs². This analysis focuses on the conceptually similar categories of value vouchers and unconditional cash transfers, for which the primary goal is to transfer value to beneficiaries rather than incentivizing certain behaviors or distributing specific goods.

Even within such programs, there are still numerous design choices for how beneficiaries are to be selected, how much they will receive, at what intervals they will receive it, and how the transfers are to be made. Some program design options may not be feasible in certain environments, but even in such cases cost efficiency analysis can help to quantify the cost implications of different design decisions.

Figure 1. Program Design Options Among Unconditional Transfer Programs



** Globally, the IRC is moving away from cash-in-envelopes to more rapidly scalable transfer methods.*

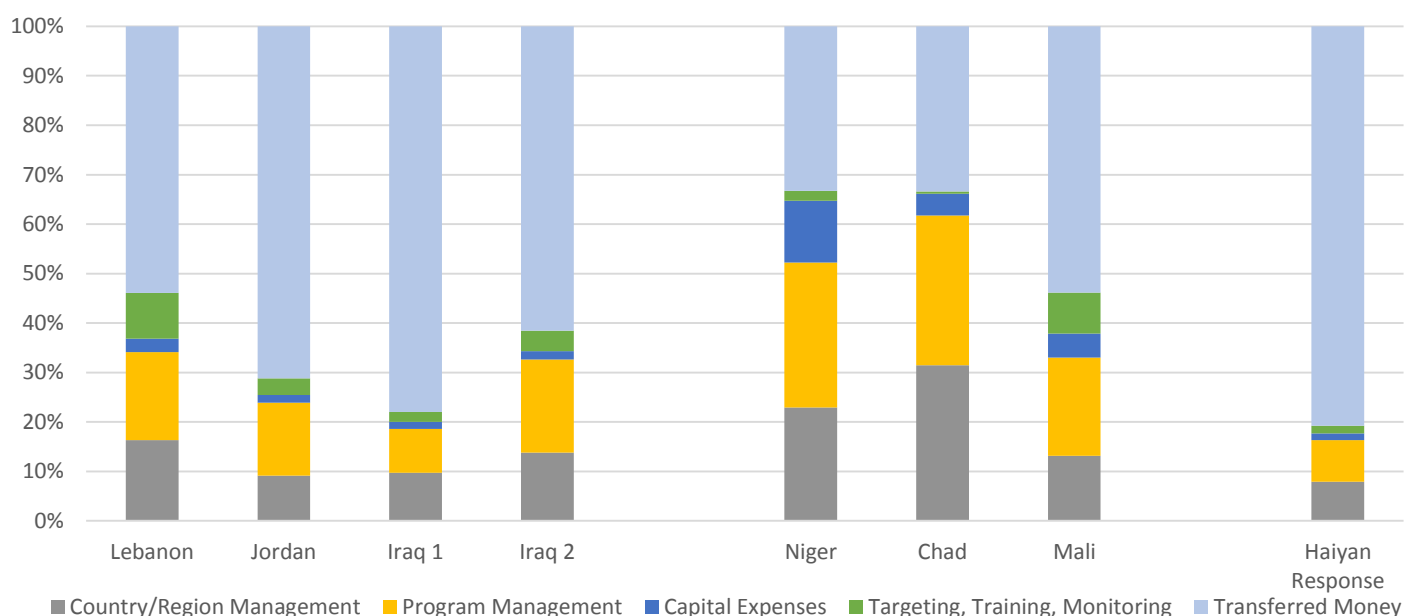
² Hidrobo, Melissa, Hoddinott, John, Peterman, Amber, Margolies, Amy, Moreira, Vanessa, Cash, Food, or Vouchers? Evidence from a Randomized Experiment in Northern Ecuador, Journal of Development Economics (2013)

Cost of Administering Unconditional Cash Programs

The IRC team has analyzed 8 cash transfer programs that were implemented in the last two years. These programs were implemented in 7 countries across 3 regions, and provided a nice mixture of different targeting and transfer mechanism.

	Country	# of Households Reached	# of Transfers Given	Value Per Transfer	Targeting Method	Transfer Mechanism
MENA	Lebanon	400	6	\$208	Vulnerability assessments	ATM Cards
	Jordan	700	6	\$170	Through women's groups	
	Iraq 1	1,064	3	\$210	Referrals from IRC staff	Hawalas
	Iraq 2	435	1	\$250		Hawalas
Sahel	Niger	500	5	\$65	Geographic + community	
	Chad	350	2	\$61	Through women's groups	Cash
	Mali	3,500	3	\$90	Community	
	Philippines	3,814	1	\$67	Geographic + community	Wire transfers

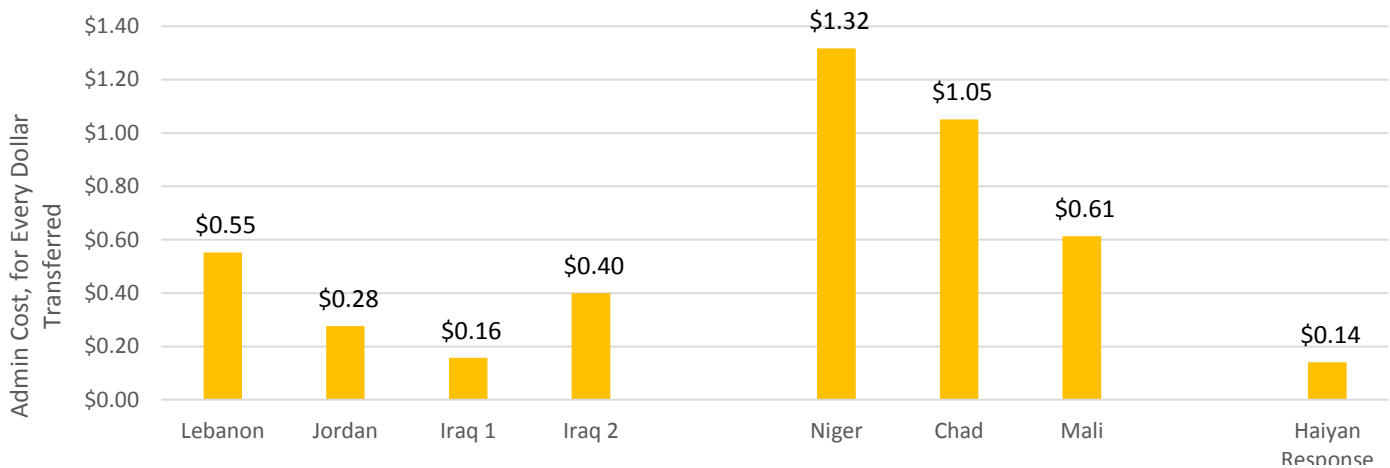
Figure 2. Breakdown of Program Spending by Category



Despite the variety of contexts and program designs, we see from the cost breakdowns that many of these programs had similar structures. Within each program, the proportion of costs dedicated to country/region support is roughly similar to the proportion dedicated to program management (In Figure 2, the yellow segment and the gray segment of each bar are usually similar in size). For countries where the local price level is lower and transfers are therefore smaller, country and program management together take up proportionally more of the total cost simply because the transfers are so cheap. In all cases, expenditures on capital goods like new vehicles or computers form only a small proportion of the total.

There are some limitations to this analysis—we have difficulty distinguishing staff time that is spent on different types of program activities, so we can't clearly separate the proportion of staff time that goes to targeting versus distribution versus monitoring. Because we do not have time-use data, we estimate the proportion of shared costs such as country director time that were spent specifically on cash distribution programming, but this means that our analysis is sensitive to that estimate. Despite these limitations, we can discern some lessons for designing more cost efficient programs.

Figure 3. Cost Efficiency of the IRC’s Unconditional Cash Transfer Programs



**Country support costs excluded, to make figures comparable to other organizations’ reported cost-transfer ratios*

At the moment, the IRC’s unconditional cash transfer programs have a wide range of cost efficiency, from a minimum of 14 cents for every dollar transferred up to \$1.32 for every dollar transferred. This demonstrates that the IRC is

capable of implementing highly cost efficient programs—for instance, GiveDirectly reports that they spend between 10 and 18 cents in administration for every dollar they transfer in stable contexts like Kenya and Uganda. At the same time, the IRC has some outlier programs which cost dramatically more per output, and do not measure up to others in the industry.

Some differences in cost efficiency are driven by contextual features that we would not or could not change. Programs run in the Sahel, for instance, transferred \$30 to \$60 per month to beneficiary households, while a program targeting Syrian refugees in Iraq transferred more than \$200 per month. This doesn’t mean we should avoid implementing programs in the Sahel: programs run in higher-cost contexts will always have lower admin cost per dollar transferred, simply because they transfer more dollars. But the wide range of cost efficiency among programs even *within* regions and countries suggests that there are also program design choices that can significantly improve cost efficiency.

Measuring Cost Efficiency

The academic literature on cash transfer programs uses the “cost-transfer ratio” to measure cost efficiency, and that is what we have reported in Figure 3. This is the ratio of all non-transfer costs, such as admin, targeting, and transfer fees, to the total value of the money that is transferred.

$$CTR = \frac{\text{Non-Transfer Costs}}{\text{Value of Money Transferred}}$$


The cost-transfer ratio is an intuitive measure, because it shows how much we have to spend on non-transfer costs for every dollar we transfer to beneficiaries. However, the cost-transfer ratio has drawbacks. Programs run in contexts where a dollar has greater purchasing power (and so fewer dollars are transferred) will always look less efficient using this metric. In Figure 3, for instance, the programs in the Sahel look generally less cost efficient than programs in the Middle East. An alternative metric that is less sensitive to differences in purchasing power is the admin cost per household per month:

$$\text{Admin Cost/HH-Month} = \frac{\text{Non-Transfer Costs}}{\# \text{ Households Served} \times \text{Program Months}}$$

This is more of a measure of operational efficiency, showing how much it cost to reach one household with transfers for one month. While programs in the Sahel look uniformly less cost efficient than programs in the Syria response region using the cost-transfer ratio, they actually have lower admin costs per household per month. The average admin cost per household per month is only \$66 in the Sahel, compared to \$74 in the Syria region.

The biggest single factor driving cost efficiency is the scale at which programs are run—reaching more households spreads the fixed costs of country support over a wider pool of beneficiaries, driving down per-household costs dramatically. The programs in this analysis reached very different numbers of beneficiaries: from 350 households in one program in Chad, to more than 3,800 households in the Typhoon Haiyan response in the Philippines. Because cash transfer programs incur fixed costs—costs that stay roughly the same no matter how many households you reach—the per-household cost drops dramatically when a program reaches more households. The IRC may want to consider guidelines that cash transfer programs should reach some minimum number of households in the future.

Figure 4. Cost Efficiency versus Scale



Program design choices about targeting method have more of an impact on cost efficiency in contexts where there is a large difference between national and international prices. The point of extensive targeting activities is to ensure that money is not misallocated to beneficiaries who do not meet vulnerability criteria, and by extension to ensure that there is enough money available for the neediest people we seek to serve. But in contexts where the price level is very low and fewer dollars are transferred to each beneficiary, the cost of giving money to a wider pool of beneficiaries and accepting some margin of error may actually be lower than the cost of extensive targeting activities, especially if those targeting activities are undertaken by international staff whose wages are dramatically higher than the value of transfers. In more expensive contexts, like Turkey or Lebanon, the cost of transfers is actually quite high and so it may be a more efficient use of resources to spend a larger portion of grant funds on targeting to ensure that the limited pool of funds goes to the intended beneficiaries.

A good approach, when considering which of two targeting schemes to use, would be to take the difference in cost between the two methods and divide this by the dollar amount of transfers you are giving. The resulting figure is the number of additional households you could serve if you used the less expensive targeting method, and you can consider whether a slightly lower margin of error among the households you do serve is worth forgoing transfers to that many recipients.

In the future, the IRC should set an internal goal of how much to spend per dollar transferred through its cash transfer programs; for example, no more than 30 to 60 cents depending on the country. These are illustrative values, intended to spark conversation rather than provide a definitive recommendation. This analysis provides some useful lessons for framing the wider conversations around benchmarks. For instance, cost-efficiency benchmarks for cash transfer programs should clearly be different for each region, so they are not driven by differences in price levels across countries. A benchmark value that excluded country management costs would be most comparable to other organization’s reported figures, while a benchmark value that included country management costs would give a more accurate picture of the resource cost of that program, so it might be worth having two versions of our benchmark. Moving forward, the Economic Recovery and Development Unit will use these insights, the results of this analysis, and input from field and technical staff to determine what benchmarks will be adopted.

The IRC’s Research & Development initiative has the stated goal of halving the cost and delivery time of cash transfer programs, and will provide operational lessons that will give country programs the tools to more easily hit these targets.

Cost Analysis at the IRC

The IRC is committed to maximizing the impact of each dollar spent to improve our clients' lives. As our CEO wrote in a recent piece in Foreign Affairs, "donors need to not just double the amount of aid directed to the places of greatest need but also undertake reforms that seek to double the productivity of aid spending." The Best Use of Resource (BUR) Initiative is focused on improving the reach and impact of the IRC by using internally available data to better understand the cost of delivering key IRC interventions. Generating evidence about cost efficiency and cost effectiveness will enable the IRC to cost and compare different approaches and their related impact, ultimately allowing decisions that achieve the best use of resources.

"Cost efficiency analysis" compares the costs of a program to the outputs it achieved (e.g. cost per latrine constructed, or cost per family provided with parental coaching), while "cost effectiveness analysis" compares the costs of a program to the outcomes it achieved (e.g. cost per diarrheal incident avoided, cost per reduction in intra-family violence). Doing cost analysis of a program requires two types of information:

- 1) Data on what a program achieved, in terms of outputs or outcomes, and
- 2) Data on how much it cost to produce that output or outcome.

Asking Ourselves, "What Did a Program Produce?"

Units across the IRC produce a wide range of outputs, from obvious items like nutrition treatment or shelter kits to more intangible things like protection monitoring or case management. Cost analysis requires us to focus in on one output (for cost efficiency) or outcome (for cost effectiveness), such as the number of items produced or the number of people provided with a service. Such outputs will not necessarily encompass all of the work that a program has done. For example, a WASH program may build water pipeline, latrines, and solid waste disposal pits, each of which could be defined as a single output. The Best Use of Resources initiative focuses on analyzing the IRC's key outputs, such as access to sanitation in refugee camps, malnutrition treatment, and case management services. Such focus is not meant to write off the other dimensions of our program's work, but narrowing in on one output allows us to compare cost efficiency across programs and contexts in a way that would not be possible if we looked at budget data at the program level. BUR staff work together with Program Quality Units to identify their most important outputs, and understand how to quantify them so that the results of our analyses are accurate and useful for programming decisions.

Asking Ourselves, "How Much Did It Cost?"

After defining the output of interest, staff builds out a list of inputs that are necessary for producing that particular output. If one thinks of a program as a recipe, the inputs are all of the "ingredients" necessary to make that recipe. Budgets contain a good deal of information about what ingredients were used, and in what quantities, but a single grant budget will frequently cover several types of outputs, or program activities across multiple sectors. This means that not all line items in a program budget will be relevant to a particular output: in order to get an accurate sense of the costs of producing the output of interest, staff categorize costs by which output they contributed to, and count only those that are relevant to the output of interest. Many of the line items in grant budgets are shared costs, like finance staff or office rent, which contribute to all of the outputs of that program. When costs are shared across multiple outputs, we have to further specify what proportion of the input was used for the specific output we are focusing on. Specifying our costs in such detail, while time consuming, is important because it provides lessons about the structure of a program's inputs. We can divide costs into categories to see whether resources are being allocated to the most important functions of program management, and to enable us to model alternative program structures and quantify the cost implications of different decisions.

This work was conducted by the Best Use of Resources Initiative at the IRC, and supported by the UK Department for International Development. Please contact Caitlin Tulloch (caitlin.tulloch@rescue.org) with any questions.

