Not all Spillover is the Same – So Don't Treat it That Way!

Ryan Bliss, Research Into Action, Inc., Portland, OR
Nicole Sage, ADM Associates Inc., St. Louis, MO
David Diebel, ADM Associates Inc., Sacramento, CA

ABSTRACT

Many survey-based spillover approaches ask respondents to provide a single, global measure of program influence on un-incented sales of efficient equipment – asking vendors and contractors to assess the program's influence on their sales or asking end-users to assess the program's influence on their purchases. Such approaches do not fully take into account that program influence may operate differently in different situations and often may operate in a manner that is invisible to any single observer. We developed an approach that assesses program influence operating in five efficient-equipment sales scenarios reflecting how the market for efficient equipment works, and identifying up to three separate pathways of program influence within each of the five scenarios. These include pathways of *indirect* program influence that can be assessed only by combining information from multiple actors.

From surveys of 33 distributors and 29 contractors that represented 50% of program savings for a large C&I program, we quantified the un-incented sales of efficient lighting equipment that occurred in each of the five scenarios. Each respondent reported the quantity of multiple lighting types sold in various scenarios, the percentage of sales that received incentives, the degree of program influence on their recommendations, and the percentage of their equipment recommendations that were accepted. For each influence pathway, we calculated program indirect influence as the product of the influence of each actor on the next one. We assessed program *direct* influence with participant and nonparticipant surveys.

In total, program influence accounted for 11.5M kWh in savings, or two-thirds of the survey respondents' total savings from un-incented lighting sales. Results showed that the program generally had greater *indirect* than *direct* influence on end-users and that vendors transmitted greater program influence than did contractors. One possible limitation of this approach is that it does not account for vendor stocking practices as a source of influence.

Potential threats to the validity of these results include the possible lack of reliability in the reported frequency with which recommendations are made and the reported influence of those recommendations; the assignment of an attribution proportion based on an influence rating scale; and restriction considering equipment recommendations as the only pathway of influence. We discuss these issues in the conclusion.

Introduction

We developed an innovative approach to estimating lighting-related spillover savings – that is, savings from un-incented sales of high-efficiency lighting equipment that resulted from direct or indirect program influence – for a utility's large C&I program. The approach built upon existing methods of assessing spillover savings in which surveyed trade allies estimate program-influenced sales of un-incented energy efficient measures. The remainder of this section describes existing approaches to assessing spillover and the limitations of those approaches and then describes the theory behind the current approach.

Definition and Assessment of Spillover

Spillover refers to reductions in energy or demand savings resulting from program influences but not arising directly from program participation (e.g., Violette and Rathbun, 2014). Spillover assessments typically distinguish between participant and nonparticipant spillover. Participant spillover generally occurs when a program participant's experience with the program leads them to install additional measures. Nonparticipant spillover occurs when program nonparticipants install measures either because of direct program influence (e.g., marketing, discussions with program staff) or because a program-influenced trade ally convinced them to carry out the upgrade.

A variety of approaches exist to assessing spillover (Violette and Rathbun, 2014; Haeri and Khawaja, 2012). Survey, or self-report, approaches are common because of the low cost. The limitations of self-report approaches are well known, and it is not the purpose of this paper to review or refute them. Because of their low cost and ease of administration, self-report approaches will continue to be used, so the objective of researchers should be to identify limitations that can be overcome with improved design and implementation.

Survey approaches to assessing participant spillover are somewhat more homogeneous than those for nonparticipant spillover. The core of the common approach is to ask program participants to report efficient equipment that they installed without program incentives or rebates but because of program influence. Most of the variation exists in how program influence is assessed and quantified and whether the approach differentiates "like" and "unlike" or "inside" and "outside" spillover (e.g., Dyson & Goldberg, 2007; Saxonis, 2007; Tetra Tech, 2011).

Greater heterogeneity exists in nonparticipant spillover approaches. One approach is to survey program nonparticipants about program-influenced but not incented upgrades (Saxonis, 2007). This approach is similar to what I have called the common participant spillover approach, except that the assessment of program influence generally focuses on awareness of the program and influence of program marketing.

The other general approach to assessing nonparticipant spillover, which seems to be growing in popularity in recent years, is to survey trade allies about their sales of program-influenced, un-incented equipment to program nonparticipants (e.g., Tetra Tech, 2011). The "core" approach here is to obtain estimates of: 1) the volume of un-incented equipment sold; 2) the amount of that equipment that was sold to program non-participants; and 2) the program's influence on those sales. Even within this core approach, methods vary in terms of how each of those estimates are obtained.

Limitations to Assessment of Spillover

It is not the purpose of this paper to review all existing methods in detail, but it is valuable to identify some limitations in some of the common existing approaches, described above, that our approach seeks to avoid.

Surveying participants or nonparticipants about program influence can reveal only *direct* program influence, but not the program's *indirect* influence acting through the actions of vendors and contractors. While they can report on the influence that the vendors or contractors had on them, their reports cannot reveal how much that influence can be attributed to the program because they do not know how much the program influenced the vendors and contractors. This limitation seems particularly troublesome for assessing nonparticipant spillover, as it is conceivable that indirect program influence on their actions is greater than direct program influence. Beyond the inability to reveal indirect influences, assessing spillover through surveys of end-users is problematic because of the low incidence of reported spillover activities, which can result in significant variability in spillover savings estimates from one sample to another (Haeri and Khawaja, 2012).

Assessing spillover through trade ally surveys can avoid the above limitations, but can have its own limitations. First, while trade allies may be able to provide reliable estimates of the percentage of sales that receive incentives, that does not necessarily translate into the percentage of sales that go to nonparticipants, as their un-incented sales may be to participants (participant spillover) as well as to nonparticipants – and there is no reason to expect that a trade ally would reliably estimate the proportion of un-incented sales that went to each group.

Further, trade ally survey methods often ask respondents to assess the program's influence on their sales of un-incented equipment. While trade allies certainly can report on the program's influence on their recommendations to their customers (e.g., Tetra Tech, 2011), that by itself does not equate to the program's total influence on the customer unless the trade ally's recommendations are completely determined by program influence. The program's influence on a trade ally's sales to customers is a function of the program's influence on the trade ally and the trade ally's influence on the customer.

The issue of assessing program indirect influence via the trade ally is complicated by the fact that it often is the case that there are at least two agents in the sales channel – an equipment vendor, such as a distributor or manufacturer representative, and an installation contractor. In such cases, the indirect program influence through that sales channel is the function of the program's influence on the vendor, the vendor's influence on the contractor, and the contractor's influence on the end-user.

A final issue is the assumption that the vendor or contractor always has some influence – or, at least, some *measurable* influence – in a sale. A vendor's stocking practices may be a more-or-less constant source of influence across sales. However, whether or not a vendor or contractor makes a specific equipment recommendation in a sale may be a variable source of influence. In cases in which the end-user specifies the equipment desired, there is no vendor or contractor influence from equipment recommendations.[1]

The failure to take account of the various roles of distributors and installation contractors introduces another concern for trade ally surveys. Any such surveys that include both equipment vendors (who sell to end-users as well as to contractors) and from contractors (who buy from vendors to sell to end-users) creates the risk of double-counting equipment that installation contractors buy from vendors and sell to end-users.

Multiple Spillover Scenarios

As the foregoing discussion shows, end-users may acquire equipment in a variety of ways. Equipment vendors sell to installation contractors but they may also sell directly to end-users, and each transaction may occur with or without an equipment recommendation. The various permutations of who sells to whom and whether or not a seller made an equipment recommendation represent differing possible pathways of program influence, as seen in Figure 1. Each arrow represents a *possible* but not *necessary* influence pathway. For example, a vendor may sell to a contractor who sells to an end-user, but without any equipment recommendations.

[1] Of course, over time the vendor or contractor's stocking practices may influence the equipment request. This certainly must be considered in assessing long-term market effects.

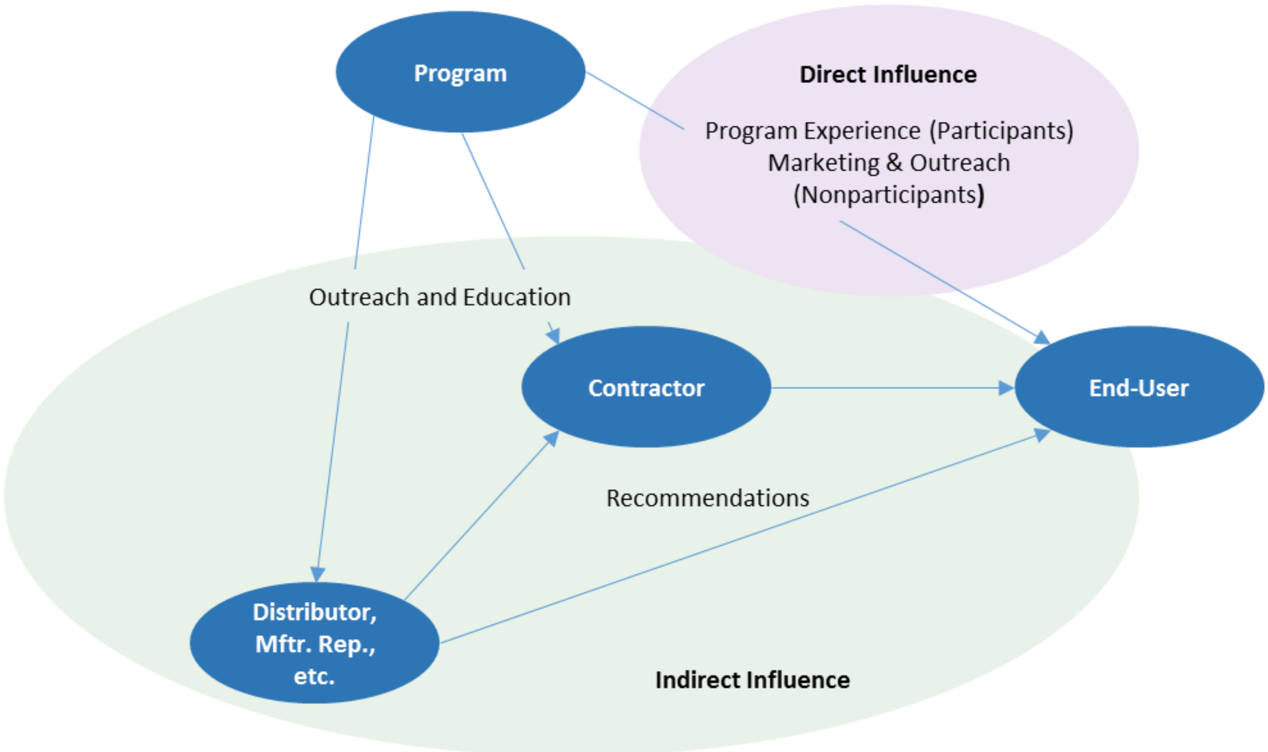

Figure 1. Pathways of program influence on end-users. Direct influence, via program experience and marketing and outreach. Indirect influence, via outreach and education to trade allies, and trade allies' recommendations to those below them in the sales channel and to end-users.

Through the above analysis, we can identify five scenarios through which equipment is sold to end-users: 1) a vendor sells directly to an end-user *without* an equipment recommendation; 2) a vendor sells directly to an end-user *with* an equipment recommendation; 3) a vendor sells to a contractor *without* a recommendation, but the contractor sells to an end-user *with* an equipment recommendation; 4) a vendor sells to a contractor *with* a recommendation, and the contractor sells to an end-user also *with* an equipment recommendation; and 5) a vendor sells to a contractor either *with or without* a recommendation but the contractor sells to an end-user *without* a recommendation. In scenario 5, the presence or absence of a vendor recommendation is irrelevant to the influence on the end-user since the contractor makes no recommendation.

The significance of identifying the multiple scenarios in which sales may occur is that, while there is always the possibility for direct program influence, the various scenarios represent different possible pathways for program indirect influence. In scenarios 1 and 5, there is no indirect program influence.[2] In scenarios 2 and 4, the only indirect program influence is via the vendor or the contractor, respectively. However, in scenario 3, program influence may be transmitted via the contractor to the end-user or via the vendor to the contractor and on to the end-user (Table 1).

[2] Other than via program influence on stocking practices. We have not considered that in the current analysis, but may attempt to incorporate it in the future.

Table 1. Five Sales Scenarios and Four Influence Pathways

Sales Scenarios*	Influence Pathways			
	1: Direct Program Influence	Indirect Program Influence		
		2: via Vendor Only	3: via Contractor Only	4: via Vendor and Contractor
1. Vendor ⇨ end-user without recommendation	✓			
2. Vendor ⇨ end-user with recommendation	✓	✓		
3. Vendor ⇨ contractor without recommendation, contractor ⇨ end-user with recommendation	✓		✓	
4. Vendor ⇨ contractor with recommendation, contractor ⇨ end-user with recommendation	✓		✓	✓
5. Vendor ⇨ contractor with or without recommendation, contractor ⇨ end-user without recommendation	✓			

* "⇨" = "sells to"

Since the program's indirect influence via any pathway is a function of the influence of each actor in that scenario upon the next actor, then the indirect influence may vary from pathway to pathway. As detailed below, we used data from surveys of vendors and contractors as well as program tracking data to estimate the total sales of un-incented high-efficiency equipment in each of the above scenarios. We used the survey data to estimate the program indirect influence via distributors and contractors and calculated the program indirect influence in each pathway. We used data from previous participant and nonparticipant surveys to estimate program direct influence on end-users.

Methods

Description of Survey

We designed separate online survey instruments for vendors and installation contractors. Both surveys asked respondents to report the number of units they sold within the utility service territory of 39 types of program-eligible high-efficiency lighting, in the following 12 categories:

- LED linear tube
- LED exterior wall pack
- LED high bay
- LED screw-in
- LED screw-in reflectors
- LED refrigerated case
- LED exit signs
- T5 high bay 150-400 watt
- T5 or T8 tube
- Ceramic metal halide
- Induction exterior fixture
- CFL screw-in

The surveys then asked questions designed to allocate the total reported sales to the five scenarios identified above. The vendor survey asked what percentage of total sales of each measure type were to contractors and what percentage were to end-users. As explained below, to avoid double-counting of spillover from sales involving both vendors and contractors, we assessed spillover savings *only* from the portion of vendor-reported sales that were to end-users.

Both surveys asked about the percentages of sales in which the respondent made equipment recommendations – the vendor survey asked this separately about contractor and end-user sales, while the contractor survey asked this only about end-user sales.

Both surveys asked respondents to report the percentage of end-user sales for which the customers reported they would apply for the program incentives, which provides an estimate of the percentage of un-incented sales.

Finally, both surveys asked respondents to rate the program's influence on their recommendations, and the contractor survey asked respondents to rate the influence of vendor recommendations on their recommendations to end-user customers; all such ratings used a scale from 1 ("no influence") to 5 ("great influence"). Both surveys assessed the respondents' influence on their end-user customers by asking what percentage of their recommendations the customers accepted.

Data Collection Methodology

The online survey approach amounted to an attempted census of lighting trade allies in the utility service territory who had done any the program projects during the 2013-2015 program cycle. We identified approximately 350 such firms in program tracking data. The tracking data included a "business type" field, allowing us to classify all members of the utility's trade ally network into vendors (those who primarily sold, but did not install, equipment) and installation contractors. We classified non-network firms based on information on the firms' websites, as confirmed in the survey. About one-third of the lighting firms were vendors and two-thirds were contractors.

We conducted the spillover surveys at the same general time as, but separately from, a process evaluation survey of trade allies conducted by telephone. We initially allocated trade allies with projects from the most recent program year to the survey frame for the process evaluation. The interviewer for the process survey asked each contact to agree to complete the online spillover survey. We sent an email invitation with a survey link to those who agreed to take the spillover survey. We also sent email invitations to complete the online survey to all lighting vendors and contractors not included in the process sample and those in the sample but not reached by the time the process survey was completed.

The email invitation to complete the online survey explained the purpose of the survey. The invitation provided contact information for key evaluation team and the utility staff. We sent up to three weekly follow-up emails to all recipients of the email survey invitation (including those process survey respondents who agreed to complete the online survey).

After three weeks in the field, we also placed calls to 40 large vendors and contractors who had not completed the survey to encourage survey completion.

The above efforts resulted in the completion of the online surveys by 33 vendors and 29 contractors. Together, those sixty-two respondents represented 50% of the lighting savings for the most recent program year.

Assigning Savings to Sold Measures

We first developed a kWh savings value for each of the 39 lighting measure categories. The kWh savings algorithm is summarized below:

$$\text{kWh Savings} = \text{Watts}_{(\text{base})} - \text{Watts}_{(\text{efficient})} / 1000 \text{ X Annual Hours of Use}$$

The baseline wattage for each set was based on commercially available nominally efficient wattages. We used the Energy Independence and Security Act (EISA) of 2007 and 2009 Department of Energy regulations to determine nominally efficient baselines. We based the efficient wattage for each lighting category on either the wattage of the actual offered measure or the midpoint wattage when a range was provided. We based annual hours of use for all interior lighting was based on the program's TRM-weighted building hours, and based exterior hours on the region's annual non-daylight hours.

Estimation of Total and Un-Incented Savings

For each survey respondent, we multiplied the number of units sold of each lighting type by the estimated per-unit savings to estimate the total energy savings from that respondent's sales of high-efficiency lighting. If a respondent reported selling a particular type of high-efficiency lighting but did not report the number of units sold, we assigned zero savings to that lighting type for that respondent.

We subtracted each respondent's incented savings from total savings to generate an estimate of un-incented savings. We had two sources for each respondent's estimate of incented savings. The first source was the estimated total savings, calculated as just described, multiplied by the estimated percentage of sales for which the customer applied for the program incentives. The second source was the program tracking data, specifically the program-incented lighting savings for projects the respondent's firm had done. To provide the most conservative spillover estimate we used the source that produced the *lower* estimate of un-incented savings for each respondent.

From the resulting estimate of un-incented savings for each survey respondent, we further subtracted any tracked spillover savings that were associated with program-incented projects ("inside spillover") that respondent's firm had done. This produced the final net un-incented sales value for each survey respondent.

Allocation of Un-Incented Savings to the Five Scenarios

For each survey respondent, we allocated the savings from the net un-incented sales to the five scenarios. We allocated vendor sales to end-users to scenarios 1 and 2 and contractor sales to scenarios 3, 4, and 5 (Table 2). The distribution of the vendor sales between scenarios 1 and 2 and of the contractor sales among scenarios 3 to 5 depended on the percentage of sales that involved recommendations.

Table 2. Allocation of Savings from Un-Incented Sales to the Five Scenarios

Scenario		Calculation of Un-Incented Sales by Scenario				
Vendor sales to end-users	1	Total un-incented sales	X	Percentage of sales in which vendor recommended equipment		
	2	Total un-incented sales	X	Percentage of sales in which vendor did not recommend equipment		
Contractor sales	3	Total un-incented sales	X	Percentage in which vendor *did not* recommend equipment	X	Percentage in which contractor recommended equipment
	4	Total un-incented sales	X	Percentage in which vendor recommended equipment	X	Percentage in which contractor recommended equipment
	5	Total un-incented sales	X	Percentage of sales in which contractor *did not* recommend equipment**		

*All contractor sales are to end-users.
**In this scenario, it does not matter whether or not the vendor recommended equipment, since the contractor did not recommend equipment, and therefore any vendor recommendations did not get passed on to the end-user.

Again, none of the scenarios includes the vendors' reported sales to contractors. That is because all vendor sales to contractors also represent contractor sales to end-users. Since this approach already counts the contractors' reported sales to end-users, adding vendor sales to contractors would double-count those sales.

Calculation of Program Indirect Influence on End-Users

We used the survey data to calculate mean program *indirect* influence through each of the three indirect influence pathways identified in Table 1, above. As Table 2 shows, program indirect influence in each pathway is a function of the influence of each actor in the pathway on the next actor in the pathway.

Table 3. Calculation of Program Indirect Influence for each Influence Pathway

Pathway	Program Indirect Influence Is a Function of…				
Vendor only	Program influence on vendor	&	Vendor influence on end-user		
Contractor only	Program influence on contractor	&	Contractor influence on end-user		
Vendor and contractor	Program influence on vendor	&	Vendor influence on contractor	&	Contractor influence on end-user

We determined that the appropriate function for combining the component influence factors is a multiplicative product of percentages, where in each case 0% is "no influence" and 100% is "complete influence." This makes sense because the product of any combination of influences cannot be greater than any of the component influences, and any component influence factor that is less than 100% will attenuate the effect of any greater other factor.

One set of influence measures already was assessed as a percentage: we assessed vendors' and contractors' influence on end-users as the percentage of their equipment recommendations that end-users accepted. As described above, contractors rated vendor influence, and both contractors and vendors rated program influence, on a scale from 1 ("no influence") to 5 ("great influence"). We converted those ratings into percentages as shown in Table 4.

Table 4. Conversion of Influence Rating to Percentages

Rating	Percentage
1	0%
2	25%
3	50%
4	75%
5	100%

Calculation of Program Direct Influence on End-Users

The current approach does not try to distinguish between un-incented sales to program participants versus nonparticipants. The approach instead uses a weighted average of the assessed program influence on energy efficiency upgrades undertaken by participants and nonparticipants from previous participant and nonparticipant surveys.

Of 488 program participants who completed the participant survey during the most recent program year, 34 reported un-incented efficiency upgrades. Those 34 respondents reported the program's influence on those upgrades on a scale from 1 ("unimportant") to 5 ("very important"). As with the vendor and contractors' influence ratings, we converted those scores to 0% to 100%.

In the most recent nonparticipant survey for the program (in the year prior to the most recent program year), 27 respondents reported on the influence of the utility's energy efficiency marketing on the decision to undertake efficiency upgrades. Again, respondents rated influence on a 1-5 scale, which the evaluators converted to scores from 0% to 100%.

Not surprisingly, the participant survey yielded a higher mean program influence score (73.4%) than did the nonparticipant survey (14.8%). To provide the weights for the two scores, we estimated the participant and nonparticipant shares of the total sales of un-incented high-efficiency equipment, using data from the vendor and contractor survey and an independent estimate of the *participant* spillover rate. The estimates used the following formulas, where x = total sales, y = participant sales, z = nonparticipant sales, q = un-incented sales, r = incented sales, and s = participant spillover rate.

(1): $x = y + z$

(2): $x = q + r$

(3): $r = y - (y * s) = y * (1 - s)$

Formulas (1) and (2) simply show that total sales are the sum of participant and nonparticipant sales, which are the sum of un-incented and incented sales. Formula (3) shows that the incented proportion of sales is equal to the total of participant sales minus the spillover (or un-incented) portion of participant sales.

We calculated the savings-weighted mean percentages of incented (r) and un-incented sales (q) from the vendor and contractor surveys, yielding values of r = .694 and q = .306.

We separately estimated a participant spillover savings rate of .015 based on the savings from un-incented equipment installed as part of incented projects (i.e., the "inside" spillover), which the program implementer tracked in the program database. Although this value likely underestimates total participant spillover, it is consistent with spillover levels found in other evaluations of nonresidential programs.[3] To the extent that it underestimates participant spillover, it produces a more conservative combined end-user influence value.

Substituting the values of r and s into Formula (3), above, and solving for y:

$$.694 = y * (1-.015) = y * .895$$

$$y = .694 / .895 = .705$$

Thus, participant sales represent 70.5%, and nonparticipant sales represent 29.5% of un-incented high-efficiency sales. We used those values with the participant and nonparticipant influence values to produce a weighted mean value for program direct influence on end-users:

$$(.734 *.705) + (.148 * .295) = .561, \text{ or } 56\%$$

Calculation of Maximum Program Influence in Each Scenario

The final stage in calculating the total spillover is to multiply the total savings from un-incented measures in each scenario by the influence value for that scenario. As Table 1 showed, however, scenarios 2, 3, and 4 each have multiple possible pathways of influence. Thus, when a vendor or contractor sells directly to an end-user with equipment recommendations (scenarios 2 and 3, respectively), there can be both program direct influence (which exists in each scenario) but also program indirect influence via the vendor's or contractor's recommendations. When a vendor sells to a contractor with recommendations and the contractor sells to an end-user with recommendations, the program may influence the end-user directly, indirectly through its influence on the contractor, or indirectly through its influence on the vendor and the vendor's influence on the contractor. Figure 3 shows the four pathways identified in Table 1.

To represent the program influence for each scenario with multiple influence pathways, we used whichever influence pathways had the greatest influence. For example, in the scenario where a vendor sells to a contractor with recommendations and the contractor sells to an end-user with recommendations (scenario 4), if pathway 3 has greater influence than pathway 1 or pathway 4, then pathway 3 defines the program influence in that scenario.[4]

[3] For example, Tetra Tech (2011), *op. cit.*

[4] The mean influence value across pathways would not be appropriate, as the various influence values in a scenario represent influence operating in separate pathways, and they do not combine in any way to determine the total or overall influence exerted in that scenario. As each is independent of the other, the one with the greatest impact is the one that represents the program influence.

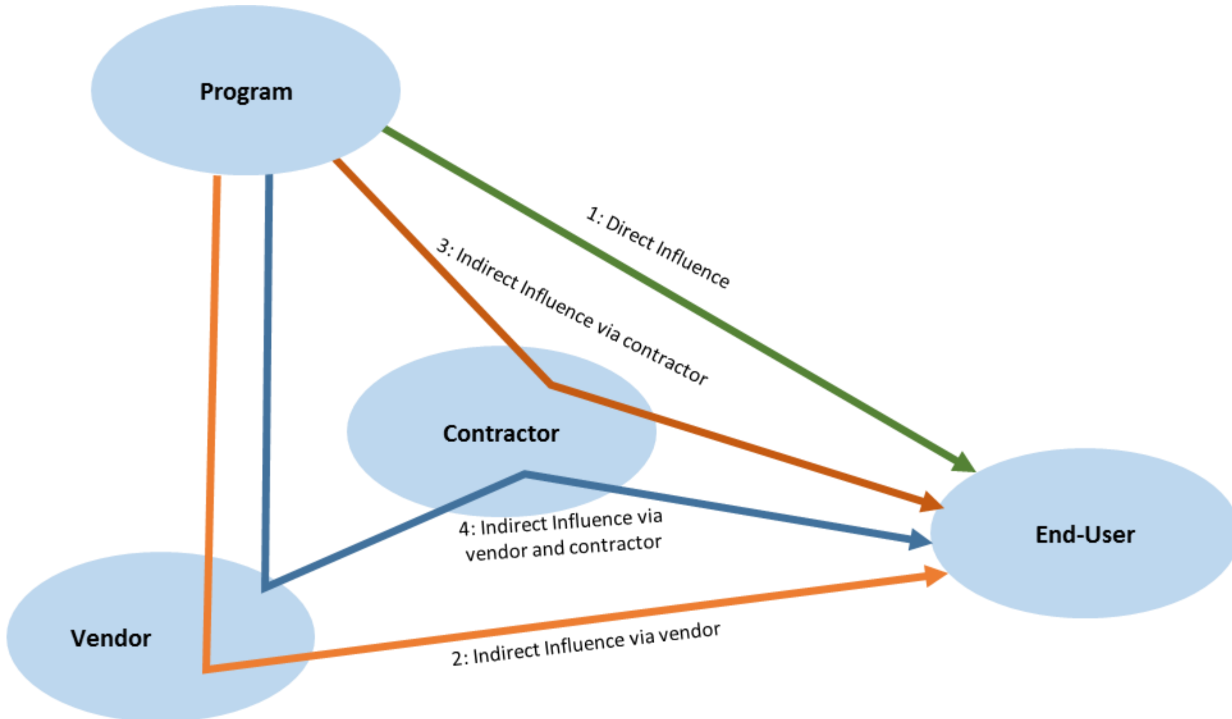

Figure 2. Pathways of program influence on end-users (revisited).

Results

Summing the spillover savings for the five scenarios produced a total spillover savings value for the surveyed vendors and contractors of 12,061,250 kWh. This represented more than 12% of the gross ex ante lighting savings for that program year. The vendor- and contractor- reported sales data were highly skewed, which, combined with relatively small samples, produced large relative errors around the mean savings values. Thus, any population estimates would have low precision. Even if we do not extrapolate to the population, the sample results themselves have some imprecision in that they relied on estimated mean influence levels. To account for this, we calculated the relative errors for the influence ratings and used them to calculate a 90% confidence interval around the spillover estimate.

Table 5 shows the results of applying our method. For each of the five scenarios, the table shows the total un-incented savings, the mean influence level for each relevant influence pathway and the maximum influence across those pathways (representing the influence within the scenario). It also shows the initial spillover estimate that results from multiplying the un-incented savings by the program influence as well as the "low" estimate that represents the lower bound of the 90% confidence interval.

Thus, we have 90% confidence that the sampled vendors and contractors represent at least 11,777,137 kWh of lighting-related spillover savings, assuming that they provided unbiased estimates of total sales and influence.

Table 5. Estimated Program Lighting Spillover – Reports by 33 Vendors and 29 Contractors

Scenario	Total Un-Incented Savings (kWh)	Mean Influence by Pathway*				Maximum Influence	Spillover	
		1	2	3	4		Initial Estimate	"Low" Estimate
1: Vendor recommends and sells to end-user	6,837,910	56%	85%			85%	5,790,329	5,409,412
2: Vendor sells to end-user without recommendation	772,213	56%				56%	434,012	410,906
3: Only contractor makes recommendations	1,815,142	56%		62%		62%	1,133,610	1,031,735
4: Vendor and contractor both make recommendation	6,775,447	56%		62%	60%	62%	4,230,275	3,851,196
5: Contractor does not make recommendation	841,848	56%				56%	473,024	447,960
Total	17,042,561						12,061,250	11,510,886

*Pathways are those shown in Figure 2:

1. Direct program influence on end-user.
2. Indirect influence via vendor (program influence on vendor X vendor influence on end-user).
3. Indirect influence via contractor (program influence on vendor X vendor influence on end-user).
4. Indirect influence via both vendor and contractor (program influence on vendor X vendor influence on contractor X contractor influence on end-user).

Conclusions

The approach described in this paper is not the only effort made to integrate information from multiple actors in the sales channel to estimate program-attributable savings. Prahl et al. (2008) laid out guidelines for integrating data from market actors and end-users in estimating net-to-gross, and Meyer (2017) recently applied such an approach to estimating free-ridership in upstream HVAC programs. Moreover, market effects research frequently attempts to integrate information from varying types of actors to assess how programs affect the market.

What is perhaps new in the approach we have described is the effort to assess the differing amounts of program influence on equipment sales in various scenarios that represent the way the market for efficient equipment works. We believe this could produce a more accurate estimate of program-influenced savings than one that relies on single average rating of program influence on sales of efficient equipment.

This research confirms that equipment vendors are important actors in the market for efficient equipment. Finally, the results demonstrate that efficiency programs may in some cases have greater *indirect* than *direct* influence on end-users, which underscores the value of maintaining strong trade ally networks to support that indirect influence.

As always, this research faces some threats to validity. First, it assumes reliability both of vendors' and contractors' reports of how often they make recommendations and of their reports of how often their recommendations are followed. In fact, we not only asked vendors how often they made recommendations to contractors, but we also asked contractors how often their vendors made recommendations to them. While the vendors, on average, reported they made equipment recommendations 68% of the time, the surveyed contractors reported, on average, that their vendors recommended equipment only 47% of the time. Similarly, we had competing sources of estimating vendor influence on contractors. One was vendors' reports of the percentage of the time that contractors accepted their recommendations – a mean of 72%. The other source, as described in the methods section above, was contractors' ratings of vendor influence – a mean of 3.5 on a 1-5 scale, which we translated as 62.5% influence. For this research, we used the contractors', not vendors', responses for estimating both frequency and level of vendor influence, both of which reduced the estimated amount of vendor-influence savings. But the greater point remains, of the need for a reliable measure.

A second potential threat is the calculation of an influence percentage from a 1-to-5 influence rating scale, such as in vendors' ratings of program influence on their equipment recommendations. For example, we translated an influence rating of 5, or "great influence," as meaning the program had 100% influence on the vendor's recommendations. The argument against this is that it implies that, without the program influence, the vendor definitely would have made different recommendations. This is a reasonable objection, which we will consider in future application of this method.

One final limitation of the approach described here is that it recognizes equipment recommendations as the only channel by which vendors and contractors deliver program influence. In particular, it does not account for vendor stocking practices or pricing as a source of influence. We will attempt to incorporate these variables into future iterations of this method.

Acknowledgements

We would like to thank all of those who provided input into the development of the approach described in this paper and who commented on the paper itself. In particular, we would like to acknowledge Greg Lovett of Ameren Missouri, Jane Colby of the Cadmus Group, Ralph Prahl of Ralph Prahl & Associates, and everyone at Research Into Action and ADM who provided feedback.

References

Dyson, C. and Goldberg, M. (2007). The Gift that Keeps Giving: A Structured Approach for Measuring Participant Spillover. Energy Program Evaluation Conference, Chicago, 2007.

Flanders, A. (2015). 101 Sources of Spillover: An Analysis of Unclaimed Savings at the Portfolio Level, IEPEC 2015, http://www.opiniondynamics.com/wp-content/uploads/2015/07/2015_IEPEC_101-Sources-of-Spillover-An-Analysis-of-Unclaimed-Savings-at-the-Portfolio-Level_Flanders1.pdf)

Haeri, H. and Khawaja, S. (2012). The Trouble With Freeriders: The Debate About Freeridership in Energy Efficiency Isn't Wrong, But It Is Wrong Headed. Public Utilities Fortnightly, 2012. (http://www.nrel.gov/extranet/ump/pdfs/troublewithfreeriders.pdf).

Meyer, J. (2017). Net-to-gross Evaluation of 2013-14 Upstream HVAC Programs (HVAC1). Prepared for the California Public Utilities Commission (CALMAC Study ID CPU0116.003), April 30, 2017.

Prahl, R., Goldberg, M., Sumi, D., Talerico, T., Tannenbaum, B., Ward, B., and Winch, R. (2008). Integrating Supply-Side Results with End-User Net-to-Gross Self-Reports. White paper prepared for the Public Service Commission of Wisconsin, March 21, 2008.

Saxonis, W. (2007). Free Ridership and Spillover: A Regulatory Dilemma. Energy Program Evaluation Conference, Chicago, 2007.

Tetra Tech (2011). Massachusetts Program Administrators Cross-Cutting C&I Free-Ridership and Spillover Methodology Study Final Report. April 18, 2011. (http://ma-eeac.org/wordpress/wp-content/uploads/Massachusetts-PAs-Cross-Cutting-CI-Free-ridership-and-Spillover-Methodology-Study.pdf)

Tetra Tech (2011). National Grid, NSTAR, Western Massachusetts Electric Company, Unitil, and Cape Light Compact 2010 Commercial and Industrial Electric Programs Free-ridership and Spillover Study: Final Report. June 23, 2011. (http://ma-eeac.org/wordpress/wp-content/uploads/Electric-PAs-Cross-Cutting-CI-Free-ridership-and-Spillover-Field-Study-Final-Report.pdf.)

Violette, D. and Rathbun, P. (2014). Estimating Net Savings: Common Practices. The Uniform Methods Project: Methods for Determining Energy Efficiency Savings for Specific Measures, Chapter 23. NREL/SR-7A40-62678. September 2014.