The Benefits and Costs of an Incentive Project or Program for State Residents: A Model for Flexible Use by Any State

This webinar is to explain a new model of incentive benefits and costs. This model is meant for use for evaluating the benefits and costs of any past or proposed incentive project, or for evaluating a past or proposed incentive program. The intended users are anyone in a state who wants to evaluate incentives: state legislative audit bureaus, state economic development agencies, and any outside group that wants to examine a state's incentives.

The webinar is accompanied by a "short" paper of 27 pages that outlines the basics of the model. If you want all the details, there also is a technical report of 208 pages, that goes into great detail on how the model works. Shortly, we also will provide on request the actual model, either as an Excel Workbook or a Python program or both; we're working on the exact Creative Commons language we will use in allowing this usage.

The model is a substantially revised version of a prior model, with the revision focused on making the model of practical use in any state, and for any incentive program. The incented project can be in any industry, and have any timetable for job creation, and the incentives can be provided by any arbitrary schedule. The user can specify any specific state and year that the project or program starts, and the model will use state-specific information to estimate the. project or program's benefits and costs.

The new model differs from some incentive evaluations in considering costs as well as benefits of incentives. The model also differs in that rather than focusing on effects on aggregate income, it focuses on effects on the income of specific people, and in particular on the per capita income of state residents. In addition to focusing on effects on average incomes of state residents, the model also reports effects by income group.

Three specific ways that the model includes costs as well as benefits are the following:

(1) Whereas some incentive evaluation models assume that 100% of incentives are decisive in inducing the incented project's jobs, this new model estimates that this percentage is much less than 100%.

(2) Whereas some incentive evaluation models assume that all increases in state personal income or output increase the per capita incomes of a state's residents, this model assumes that much job growth simply increases population growth, with little effect on per capita income.

(3) Whereas some incentive evaluation models ignore the cost of paying for incentives, the new model allows for the budget costs of incentives to have some adverse effects on a state's economy.

Now, more elaboration on those three points.

On the first point, some evaluations of incentives treat the project as if "but for" the incentives, the project would not have been undertaken. But because incentives are often only a modest percentage of total costs, the new model projects that for many incented projects, the "but for" will be less than 25% -- that is, 25% of less of the projects would not have occurred in the state "but for" the incentives.

In addition, the model allows you to be agnostic about "but for", and simply impose a "but for", or to determine what "but for" is needed for the incented project to have net benefits for state residents. This avoids the highly implausible assumption that incentives are always decisive.

In operation, the model seeks to estimate the expected benefits of an incented project. So, the project's jobs are multiplied by the "but for" percentage to get expected jobs, which in turn yields expected benefits.

On the second point, studies of local labor markets suggest that on average, when an incented project creates new jobs in a state economy, that after 7 years or so, on average only 20% of the new jobs go to local residents – the other 80% go to in-migrants.

This seems odd. Surely new jobs hire locally more than 20%. But what this overlooks is job vacancy chains. When new jobs are created locally, they are filled in 1 of three ways: hiring local residents who are already employed, hiring local residents who are non-employed, and hiring in-migrants. When already-employed local residents are hired, this creates a job vacancy that is filled in the same three ways. These job vacancy chains are only terminated when every new job created either results in a non-employed local resident getting a job, or an in-migrant coming in for a job. As a result, the in-migration effect of new jobs is much greater than the initial hiring might suggest.

This population increase has a number of consequences, including that new jobs in a state increase state personal income per capita a lot less than they increase state personal income. In addition, greater population increases local housing prices, which helps some local residents who own property, but who tend to disproportionately have higher incomes, while hurting some others on fixed incomes.

On the third point, incentives will increase state and local tax revenue, but they are highly unlikely to pay for themselves. One important reason for this is that because job growth also increases population growth, the increased population will increase demand for state and local public services to keep public service quality constant, for example to hire new K-12 teachers so class sizes can be kept the same as population and student numbers go up. More realistic models find that fiscal benefits offset less than one-quarter of incentive costs.

The net fiscal costs of incentives have economic costs for a state's residents, in terms of higher taxes and local public spending, and the economic effects of higher taxes and lower public spending. In particular, lower spending on K-12, by lowering school quality, will have adverse effects on future earnings of K-12 students.

So, that's some of the logic of the model – it accounts for some of the costs as well as benefits that occur when we decide to incent a particular project or have an incentive program. How is this model actually implemented.

For a particular incentive project, at a minimum, the user has to specify the following:

- The number of incented jobs, and how they evolve over time.
- Dollars of incentives provided each year.
- What industry it's in
- The multiplier for the project, which you can get from BEA-RIMs, or IMPLAN, or which, based on prior research, is likely to be around 2.
- The state and starting year of the project, which brings in state and time period specific information to do the model's simulations.

To illustrate, I consider a hypothetical simulation program. I imagine that this project is in Michigan, starting in 2024. I assume the project is in an "average" industry, and has a fairly typical multiplier, of around 2. I assume the project has 10,000 jobs, and that the incentives are $5,000 per job for 15 years, but paid with a one-year delay.

One aspect of the model that the user also has to specify, which may be sometimes challenging, is what to assume about what happens to the incented jobs after the incentive period is over. The model simulates economic effects on a state's residents for 80 years, so we need to know what happens over an 80-year period. Incented jobs may expand further after the incentive period, or may contract or close down. And if the project closes down, some other firm may take over that website and offer similar numbers of jobs. My own advice is that a relatively neutral assumption is to assume incented jobs persist at the same level, neither expanding nor contracting. This is consistent with how local economies tend to respond to job shocks.

The model then produces many outputs. But one of the most important outputs is the estimates it provides of the effects on the present value of per capita income of state residents. This includes effects on the real value of all types of income for state residents. It also includes a breakdown by "income quintile". In such an income quintile analysis, we rank all households by household income, adjusted for household size, and then divide into five groups with equal population in each group.

For presentation purposes in this webinar, I will aggregate into just 4 types of income, and only look at effect on 3 of the 5 income quintiles: the lowest income quintile, the middle-income quintile, and the highest income quintile.

So, if every quintile had the same income, each quintile would have 20% of total income. The lowest income only has 5% of total income, so its average income is one fourth the average. On the other hand, the highest income quintile has about half of total income, so its average income per capita is about 2.5 times the average. Overall, the lowest income quintile has about one-tenth of the per capita income of the highest income quintile.

For this hypothetical project, we project the present value of increases in per capita income of 4 types of income. First, the project, by creating jobs, has some labor market effects in increasing earnings per capita, principally by increasing employment to population ratios. In addition, by increasing population, the project will increase property values, which produces capital gains for local residents who own property. But paying for incentives has some costs, both directly in the form immediately of higher taxes and lower public spending. But it also has some indirect effects, by reducing future earnings due to lower K-12 spending and quality.

Net effect is $200 million net benefit for state residents, accounting for all incentive benefits and costs. The labor market benefits and property value benefits outweigh the direct and indirect costs of paying for the incentives.

But not all income groups benefit. In particular, note that the lowest income quintile loses. Why? Well, the lowest income group does get more than its share of the labor market benefits. Its initial income share is about 5%, and 86 out of 738 is more than 10%, so it gets more than double its current share. If only labor market benefits mattered, the lowest income quintile would make out great, by getting much-needed jobs.

However, the lowest income quintile gets very little of the property value benefits. In addition, it pays more than its share of the incentive costs, as state and local tax increases tend to go disproportionately on lower income groups. Most importantly, the long-run costs of cutbacks in education quality puts a disproportionate burden on the lowest income group, as the evidence suggests that children in all income quintiles have similar long-run earnings benefits from public schools, and the lowest income quintile tends to have slightly higher numbers of children per households.

In contrast, the middle-income gets more than its share of the benefits of this incentive project. Its baseline income share is 14 percent, yet its gains are almost half of the income gains overall. Why is that? Well, the middle-income group gains a large share of the labor market benefits as well. And although it loses from paying for incentives, both directly and indirectly, those losses don't loom quite as relatively large.

The upper income quintile gains, but slightly less than its share of income. Its baseline income share is slightly over 50%, but its share of net benefits is less than one-half. The highest income quintile gets a disproportionate share of the property value benefits and suffers relatively less from paying for incentives, but tends to get less than its baseline share of labor market benefits – this quintile already is pretty much fully employed in good jobs.

Now, one concern I have heard about this model is that the model does simulations of incentive benefits and costs 80 years into the future. The argument against that is simple: who believes projections 80 years into the future?

One response is to note that the model produces results by year, so it is quite possible for a user to project effects only for whatever period is desired, for example 10 years.

But a better response is that the model does not really seek to project the entire economy into the future. Rather, it simply asks how the world will be different because of the incentive project, not just 10 years from now, but even decades later. And if we restrict our attention to only the first 10 years, for example, we ARE making an assumption about the effects of this incentive program.

Specifically, we are assuming that after 10 years, the incentive program will have zero effects. And in many cases, that is not a reasonable assumption.

To illustrate this, for this hypothetical incentive program, we can look at net benefits over time, and at cumulative net present value of net benefits as of various time period. The latter reflects the sum, using discount rates, of the former measure, up to the last year.

If one looks at cumulative present value, one would get a misleading picture of the net benefits of this incented project if one truncated things before 60 years. The cumulative present value goes up for 6 years, then goes down again until year 16, then goes up again until year 25, and then goes down again over the next 35 years.

What is going on here? Initially, the project creates jobs, which creates some labor market and property value benefits over the first six years. But then the incentive costs, which continue until year 16, begin to have direct costs for state taxpayers, and indirect economic costs for state residents. Also, the initial labor market benefits begin to fade, and the higher property values lead to higher property taxes and some reduction in real incomes of people on fixed incomes who are renters. At year 17, the incentive costs end, and this both benefits state taxpayers, and also provides a boost to the economy. But this fades somewhat over time. And starting around year 26, we begin to see sizable costs from the reduced education spending due to the incentive costs. These costs of lower K-12 spending continue until about year 60, when most of the people who were in school during the incentives are through with most of their careers.

To put it simply: ignoring long-run effects of incentives is assuming that no changes in state and local budgets have long-run productivity effects. And that is often a bad assumption.

As mentioned, one aspect of the model is that it emphasizes that the effects of incentives don't only depend on the industry of the project and how the incentive is timed. For example, benefits also depend on where in the state the project is located, and how the project is financed.

For example, suppose we take the same project – same 10,000 jobs, same $5,000 per job per year for 15 years, same multiplier – but assume the local economy has a much lower employment to population ratio or employment rate. So, since I am imagining a Michigan project, suppose that the project, rather than being located in an "average" Michigan county, is located in a more economically distressed place in Michigan such as Flint, Michigan, or some distressed rural county. In that case, the benefits will be much higher because there are more available local non-employed persons to hire, and thus where will be less in-migration. And the benefits will be distributed more progressively.

The next table shows this. Because of the lower employment rate, labor market benefits are greater – the employment rate goes up more, and the population goes up less. As a result, not only are labor market benefits higher, but fiscal benefits are higher – not as much increased population has to be accommodated. And education cutbacks are less extensive.

Overall, net benefits almost triple, from $200 million to $599 million. And the lowest income group now gain. In fact, its gains of $47 million, out of the total gain over all quintiles of $599 million, are more than its baseline income share of 5 percent. This occurs because this lowest income quintile particularly depends on the labor market benefits, and also benefits more when incentives have fewer fiscal costs.

Net benefits also depend on how incentives are financed. For example, if we return to the baseline model, with the Michigan average employment rate – we're no longer in a distressed area. This baseline model assumes that 50% of the incentive is financed by tax increases and 50% by public spending cuts. Of those public spending cuts, 20% is assumed to come from K-12 public spending. This percentage is based on Michigan averages.

Suppose we instead assume that 100% of the incentive is financed by public spending cuts, with the same 20% of public spending coming out of K-12. This doubles the effects of paying for incentives on K-12 spending and school quality.

As this simulation shows, under these assumptions, the overall project has much more negative effects due to how education cutbacks reduce future earnings. In addition, because of lower future earnings, there are lower future tax revenues, resulting in lower fiscal benefits and greater net fiscal costs. And the lower earnings also slightly reduce demand and therefore jobs in the state.

The result is that the project switches from having $200 million in net benefits, to $215 million in net costs. This is mostly due to the greater education cutbacks.

These greater education cutbacks particularly hurt the lowest income quintile. In the baseline scenario, the lowest income quintile does not really gain or lose much. But in this new scenario, the lowest income quintile suffers losses in income that at $98 million almost half the overall losses for all quintiles.

In sum, this model:

- Can be used for any state and any incentive program and is attuned to the characteristics of the state and of the program.
- Shows how incentives affect the economic well-being of state residents in various ways and allows results to be disaggregated by effects on different types of income, including fiscal benefits, as well as being disaggregated by income group.
- Can be adapted to incorporate various assumptions. For example, the model can be adapted to show how effects vary based on various assumptions about "but for", as well as assumptions about where in the state the project is located, or how the project's incentives are financed.