ePolicy

Engineering the Policy-making Life Cycle

Seventh Framework Programme – Grant Agreement 288147

Deliverable 5.5
Second demonstration of the testbed

Document type:	Deliverable
Dissemination Level:	Public
Editor:	Lars Kotthoff
Document Version:	1.0
Contributing Partners:	UCC, US
Contributing WPs:	WP5
Estimated P/M (if applicable):	3
Date of Completion:	31 Jul 2014
Date of Delivery to EC	31 Jul 2014
Number of pages:	21

ABSTRACT

This report demonstrates the testbed for formulating and exploring complex policy design problems motivated by the case-studies settings provided by the consortium, specifically the allocation of grants to citizens in the Emilia-Romagna region of Italy and Bologna, its capital, to incentivize the installation of photovoltaic systems. We explore a number of different parameter settings and scenarios and investigate their effect on the installed power capacity and cost of the individual policy instruments. We draw conclusions as to the most effective incentive instrument and investigate an alternative, potentially more optimal, way of allocating funds to bids.

Authors of this document:

Lars Kotthoff, Yulia Malitskaia, Barry O'Sullivan, Helmut Simonis, Nic Wilson, Peter George Johnson

INSIGHT Centre for Data Analytics
University College Cork, Ireland
email: larsko@4c.ucc.ie, y.malitskaia@4c.ucc.ie, b.osullivan@cs.ucc.ie, h.simonis@4c.ucc.ie, n.wilson@4c.ucc.ie
and
Centre for Research in Social Simulation
University of Surrey, UK
email: peter.johnson@surrey.ac.uk

Contents

This page has been intentionally left blank.

1 Executive summary

This deliverable describes and demonstrates the results the integrated social simulator and incentive design component achieves on the ePolicy case study of the Emilia-Romagna region in Italy and its capital, Bologna. We evaluate the four policy instruments (interest rate support, grants, feed-in tariffs, and tax benefits) considered for incentivizing photovoltaic solar panel (PV) adoption by households. We analyse the effect of regional policies on household PV adoption; specifically the effects of the regional incentive budget (in Euro) and the type of incentives used (both regional and national) on the power produced in the region and the costs to regional government.

The main changes compared to the first demonstration of the testbed in deliverable 5.4 is a more detailed evaluation of the effect of the different types of incentives with a larger number of different scenarios, the investigation of the counter-intuitive behaviour where the total capacity decreases despite an increase in budget described in deliverable 5.4, and a preliminary evaluation of a different means of computing an optimal allocation of funds to bids.

In addition, the way the component is integrated into the overall ePolicy system has changed. In particular, the incentive design and social simulator are run earlier in the process to determine the maximum energy that can be produced with a given budget in a given scenario. This output is then used in the other parts of the system to inform the decisions made there. For the incentive design/social simulator, this means that only the optimisation objective that maximises the power is relevant and the target energy parameter becomes superfluous. There are fewer parameters to set and the component is easier to use.

We present comprehensive experiments that explore a range of different parameters and a total of 3584 different scenarios that consider not only the Emilia-Romagna region as in previous deliverables, but also Bologna on its own. In particular, we investigate a wide range of regional budgets, budget distributions and combinations of incentives. Our experimental results demonstrate that regional incentives play only a minor role compared to the national incentives, although they seem to act as a catalyst. As the budget for regional incentives increases, the total installed capacity does as well in general if national incentives are present.

The most efficient incentive mechanism in terms of the ratio of cost to added installed capacity appears to be tax benefits. While it does not quite achieve the level of added capacity feed-in tariffs do, the cost it incurs is much smaller.

Preliminary results we present for an alternative scheme to allocate funds to bids that aims to optimise the allocation of funds to bids over the entire simulation rather than each time step individually are very promising. While further research is required to make this approach suitable for inclusion in a decision support system such as the ePolicy system, we demonstrate its potential.

2 Introduction

The aim of the integrated social simulator and incentive design component is to allow policy makers to evaluate the effects their decisions are going to have with respect to individuals installing solar power. In conjunction with the other components provided by the ePolicy project, it allows to design and evaluate energy plans.

The population of the case study region of the ePolicy project, Emilia-Romagna in Italy and its capital, Bologna, is modelled as a social simulation with simulation agents representing specific parts of the population determined by the geography and demographics. The policy maker specifies what form the support of the regional government for individuals who wish to install solar power takes and the budget available for it. Agents may make the decision to bid for government support, where a bid includes the amount of funding requested and the power capacity of the proposed installation.

The social simulation runs in time steps of one year. At the end of each time step, the bids the agents generate are collected. The incentive design part uses optimisation technology to determine, given the budget and other parameters the policy maker has specified, which bids should be funded. It thus ensures an optimal allocation of the available funds with respect to the policy maker's objective.

In this deliverable, we present and evaluate the results across a range of different funding scenarios, with different budgets, budget distributions and other parameters. This demonstration shows a comprehensive overview of the results that can be achieved with this component and showcases its capabilities.

The remainder of this document is structured as follows. First, we briefly remind the reader of the architecture of the integrated social simulator and incentive design component in Section 3. Section 4 describes the scenarios and parameters used in the demonstration and experimental evaluation of the integrated component. After describing and discussing the results in Section 5, we explore the effect observed in deliverable 5.4 where the total capacity decreases despite increasing budget in Section 6. In Section 7, we present preliminary results for a novel allocation approach that attempts to optimise the allocation of the incentive budget over the entire duration of the simulation rather than for individual time steps only. We conclude in Section 8.

3 Integrated social simulator and incentive design component

The integrated social simulator and incentive design component has been described in deliverable 5.3 [6]. The high-level interactions between user, optimisation and simulation are depicted in Figure 1 as a reminder. A description and some preliminary experimental results with particular focus on the use of optimisation technology in the context of agent-based simulation has been published in [5].

The main changes compared to the previous version of the integrated social simulator and incentive design component (deliverable 5.4 [7]) are as follows.

- The optimisation objective and target energy parameters have been dropped because of the changed interaction with the rest of the ePolicy system.

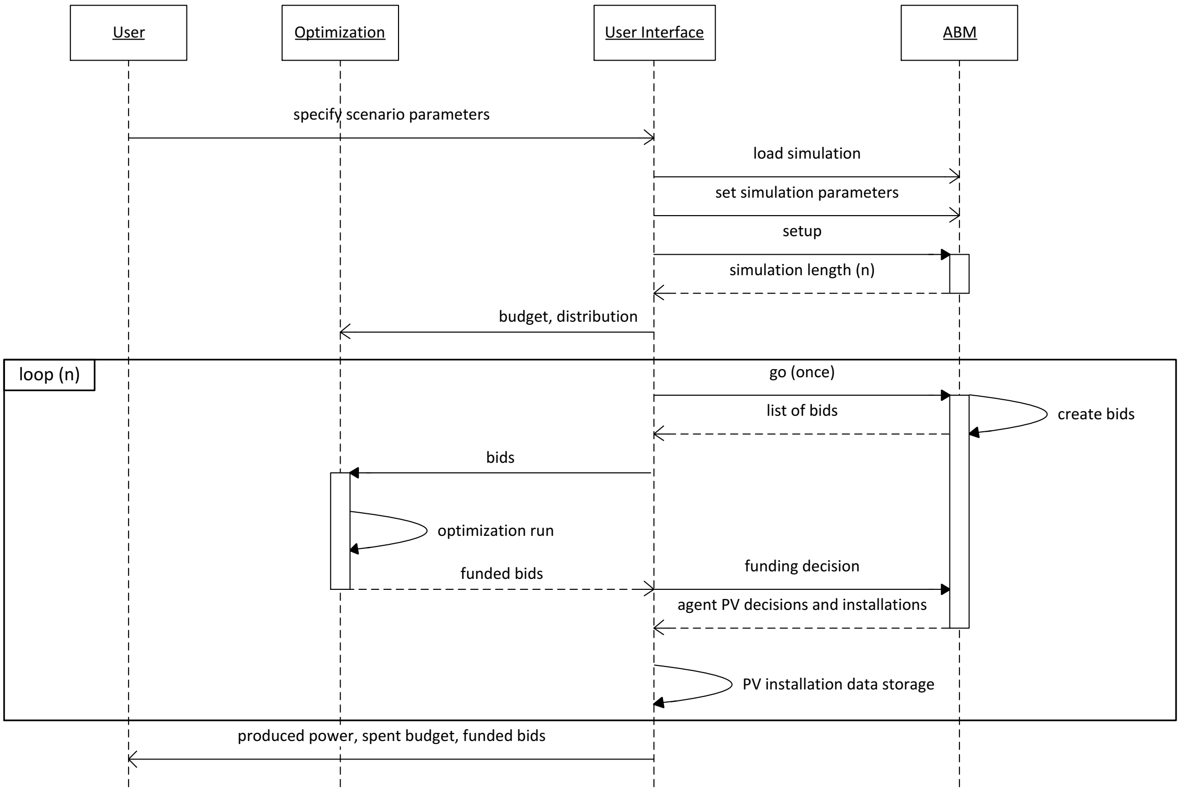

Figure 1: Sequence diagram of the integration of the social simulator and incentive design.

- Additional verification and validation has been performed.

We retain the parameters to set the optimisation objective and target energy in the stand-alone component to allow more flexible use.

3.1 Installation

The integrated social simulator and incentive design component is available at `http://4c.ucc.ie/~larsko/downloads/pver-0.2.zip`. The software can be run as follows. Extract the deliverable file `pver-0.2.zip` to a location of your choice. Inside the extracted folder, run the following commands on Linux.

```
export LD_LIBRARY_PATH=lib
java -Xmx2G -Djava.library.path=lib -jar pver-0.2.jar
```

On Mac OS X, the following commands should be run.

```
export DYLD_LIBRARY_PATH=lib
java -Xmx2G -Djava.library.path=lib -jar pver-0.2.jar
```

This will start the application. For more information on its usage, refer to Deliverable 5.3 [6].

4 Experimental parameters

In total, we considered 3584 different scenarios for the demonstration. The simulation for each scenario was repeated five times, as determined in deliverable 5.4 [7]. The scenarios evaluate the effects of the interactions of the key input parameters and the various levels they can be set at.

The changed interaction with the rest of the ePolicy system means that only the optimisation objective where the energy output is maximised is considered and the target energy becomes superfluous. Both parameters are therefore fixed and dropped from the input. They do not appear in the list of surveyed parameters anymore.

In detail, the parameters we considered were the following.

Region Emilia-Romagna and Bologna

FIT on and off

Tax relief on and off

Grants on and off

Interest rate support on and off

Budget Distribution first come first served, ramp-up, ramp-down, and even

Regional Budget €1,000,000, €2,000,000, €3,000,000, €4,000,000, €5,000,000, €6,000,000, €7,000,000, €8,000,000, €9,000,000, €10,000,000, €20,000,000, €30,000,000, €40,000,000, €50,000,000, €60,000,000, €70,000,000, €80,000,000, €90,000,000, €100,000,000, €200,000,000, €300,000,000, €400,000,000, €500,000,000, €600,000,000, €700,000,000, €800,000,000, €900,000,000, and €1,000,000,000

We evaluate the full cross-product of all the above parameter settings. Not all of the particular values we explore are necessarily realistic, but they demonstrate that the integrated component is able to deal with a large range of inputs and reliably deliver outputs.

Compared to the evaluation in deliverable 5.4 [7], we consider Bologna as a region in addition to Emilia-Romagna, we change the national as well as the regional incentives, and consider more different budgets.

5 Experimental results

We average the results achieved at the end of each simulation over the five repetitions for each scenario and start by analysing the relationship between the total capacity installed as a result of the scheme and the total cost for the incentives. The total cost includes the cost for the regional incentives and the cost for the national incentives. Figure 2 shows this relationship for the Emilia-Romagna region. Each point in the graph is an individual scenario. The colour of the point denotes what kind of incentives were used.

The experimental scenarios separate into three "clusters" that are determined by the type of incentive used. The dots with orange/brown colours can all be found at the bottom left of the graph and denote the scenarios where none or only regional incentives were used. The dots with blue colours at the top left of the graph denote the scenarios where tax relief national incentives were used, with and without regional incentives. The third cluster of points, green and purple/pink colour at the top right of the graph, denotes the scenarios that used national feed-in tariff incentives, with and without the other three types of incentives.

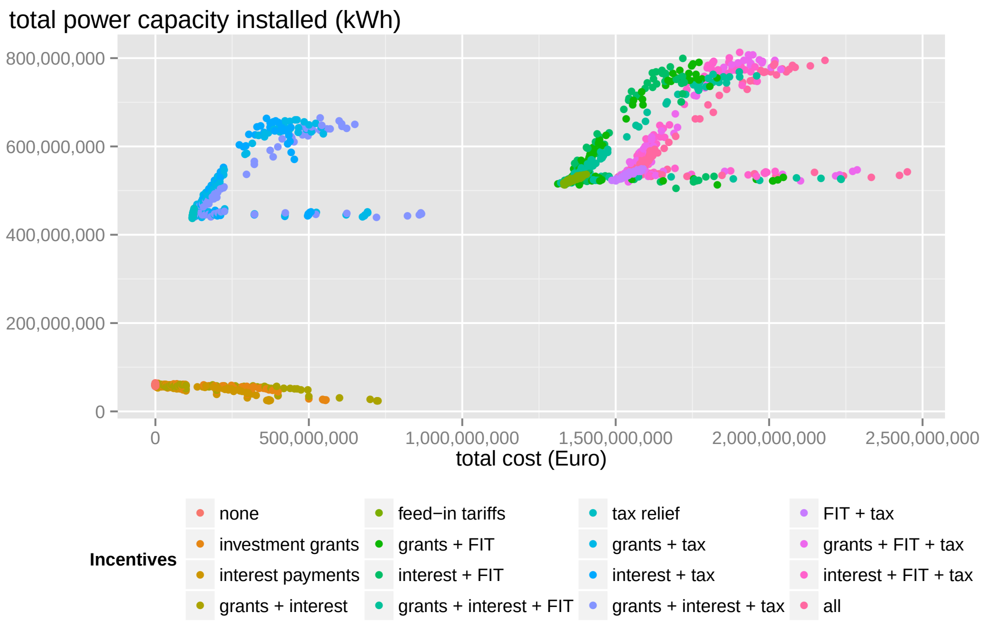

Figure 2: Total power capacity installed over total cost for the Emilia-Romagna region.

The general shape of the graph corresponds to what one would intuitively expect – as the installed capacity increases, so does the cost. More installations require more funding, driving up the cost. The graph also demonstrates the impact of the different incentive instruments. While the regional incentives have a relatively low impact, the national ones (feed-in tariffs and tax benefits) make a huge difference, increasing the installed capacity by up to an order of magnitude when switched on. The regional incentives become actually less effective as the cost increases if no national incentives are present (orange/brown dots).

The overall highest output is achieved when feed-in tariffs are used as incentive instrument, with the addition of tax benefits increasing the cost significantly, but not the total capacity (green vs. purple/pink dots). Tax benefits combined with regional incentives (blue dots) achieve only slightly reduced capacity, but at much lower cost.

Figure 3 shows the same relationship for the Bologna region. The figure is very similar to Figure 2, showing the same patterns. However, both installed capacity and total cost are much lower, as expected. In addition, the installed capacity seems to decrease as the budget increases beyond a certain point – the same phenomenon we observed for the Emilia-Romagna region at higher budget levels in deliverable 5.4.

When examining the relationship between the installed capacity and the budget for regional incentives, the same general pattern can be seen (cf. Figure 4). The capacity due to national incentives is much higher than that due to regional ones. As the budget increases, so does the installed capacity, but only up to a certain point. When national incentives are used, the total capacity starts to decrease as 1 Billion Euro is reached for the regional budget. If no national incentives are used, the installed capacity does not increase with increasing budget.

Curiously, the budget for the regional incentives has a much bigger effect when national incentives are used as well. This suggests a synergistic effect between national and regional

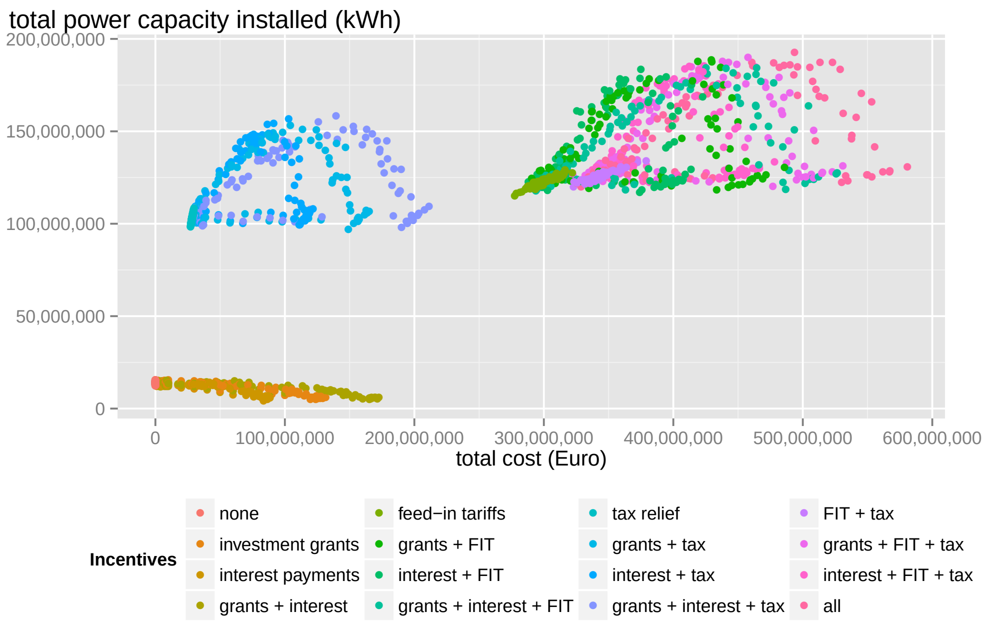

Figure 3: Total power capacity installed over total cost for the Bologna region.

incentives – increasing the regional budget stimulates demand, but achieves better results only when national incentives are present as well. The graph for the Bologna region looks again very similar and is omitted here.

5.1 Effectiveness of incentives

Figure 5 sheds more light on the effectiveness of the different incentives. It becomes very clear that regional incentives alone achieve much less than when combined with national incentives, or even national incentives on their own. Again feed-in tariffs achieve the overall best result, with added tax incentives achieving only a slight improvement and the graph for Bologna shows the same picture.

The relationship between total cost and type of incentives used is further explored in Figure 6. Feed-in tariffs are clearly the incentives instrument that incurs the overall highest cost, which explains why they are the most effective one – people participating in the scheme get most money through it. In fact, the median cost for scenarios that do not have feed-in tariffs is almost an order of magnitude lower than for those including feed-in tariffs. The same is not true however for the total capacity installed.

Based on our experimental evaluation, we conclude that for the scenarios we have considered, tax benefits are the most efficient incentive mechanism based on their cost to benefit ratio. While they are much cheaper than feed-in tariffs, they achieve almost the same level of PV adoption in terms of capacity installed.

The regional incentives play only a minor role, although our results suggest that they act as a catalyst for national incentives.

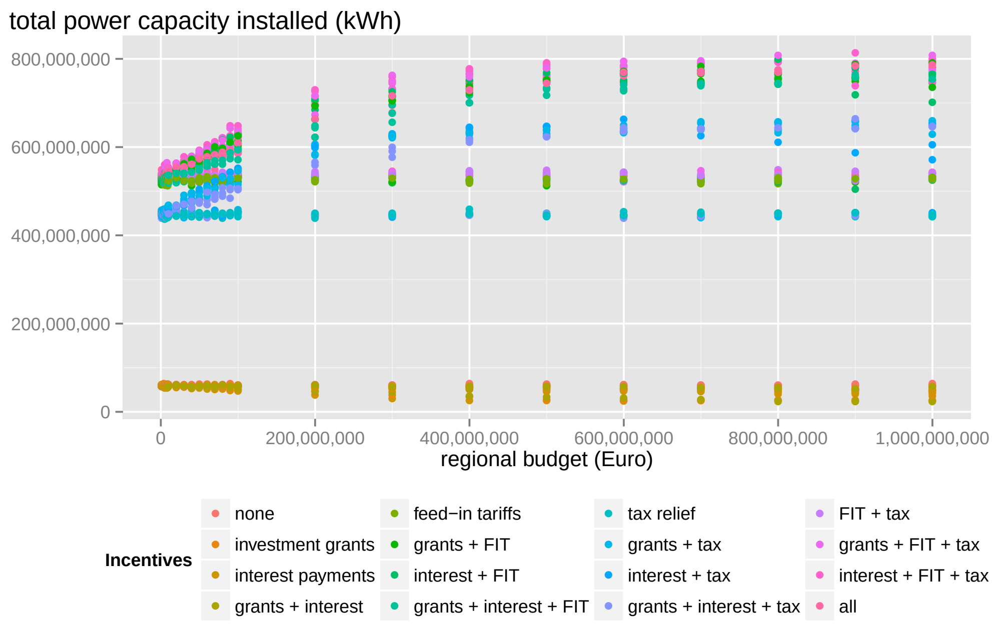

Figure 4: Total power capacity installed over budget for regional incentives for the Emilia-Romagna region.

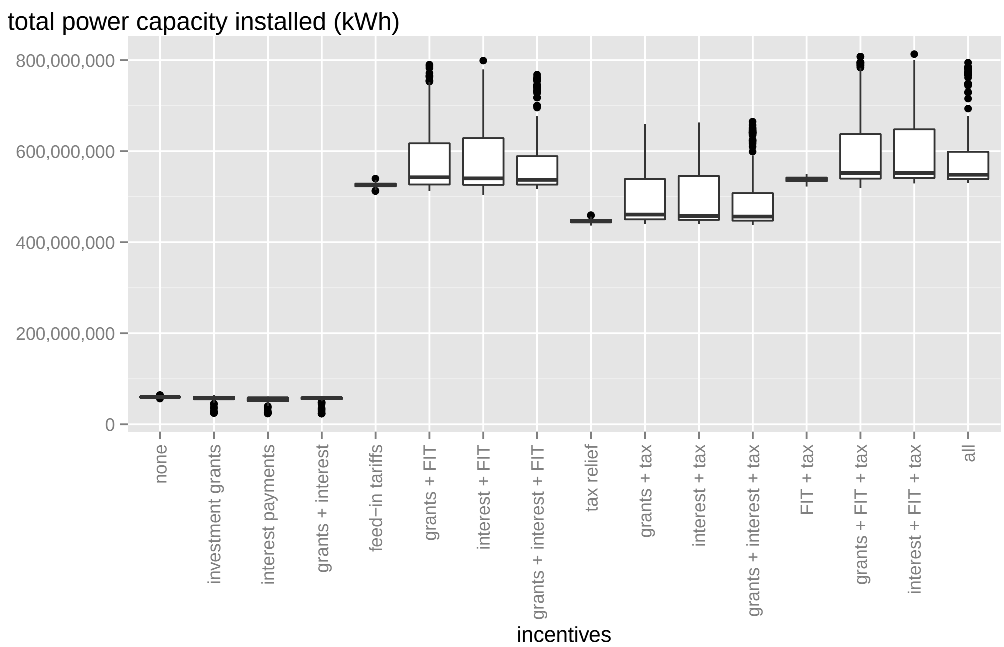

Figure 5: Total power capacity installed over types of incentives used for the Emilia-Romagna region.

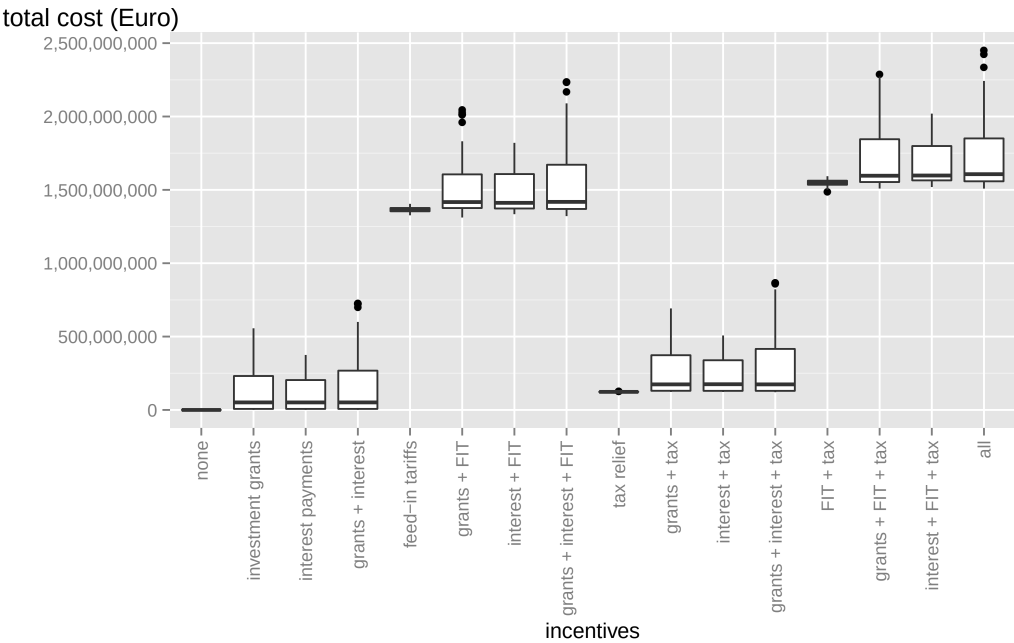

Figure 6: Total cost over types of incentives used for the Emilia-Romagna region.

6 Investigating of decrease in capacity despite increasing budget

As briefly described in deliverable 5.4, we observed in our experiments that when the budget for regional incentives increases beyond a certain point, the total capacity decreases. This is also present in the experimental results presented in Section 5. Our intuition was that this is caused by effects in the simulator that start to occur when the level of funding is much higher than the level of demand.

We investigated this issue in close cooperation with the University of Surrey, who designed and implemented the social simulator. Unfortunately, despite extensive experiments we have so far been unable to establish the source of this issue beyond doubt. The main problem with our investigation is the sheer scale of the problem – the simulation involves many agents and it is infeasible to track each and every one of them throughout the time steps of the simulation.

Below we offer some intuition along with evidence for the cause of this behaviour. We will continue to investigate and present any additional findings in future deliverables; however, the behaviour only occurs at very large budgets (1 Billion Euro and beyond), which are not realistic. While this investigation may be interesting and worthwhile from a scientific point of view, it is therefore largely irrelevant for the ePolicy project. We are therefore planning to minimise the resources spent on this tangential issue and instead focus on the dissemination of the project and the involvement of policy makers.

What we observed throughout all the experiments run for this investigation is that the share of national incentives is very high for small regional budgets and decreases as the regional budget increases. That is, the more regional budget is available, the more people are taking advantage of this rather than only the national incentives. For very large budgets, all installations are supported by both national and regional incentives.

We also observed that the drop in installed capacity with increasing budget does not occur when the budget distribution of the time steps is first come, first serve – that is, the entire (remaining) budget is available in each time step. Our intuition is that for the other budget distributions for small budgets, very few people get accepted for regional incentives. Many of those that do not get regional support get just national support. In subsequent time steps, the vast majority of households do still not receive regional incentives, which are now also unable to get support from national incentives. This explains the effect where the PV adoption stagnates despite increased regional budget.

For higher regional budgets, most applicants receive regional incentives and go on to receive national incentives as well. However, there are still some households that get only national support, then in later rounds other households get regional support. Thus we see a more constant growth. For much higher regional budgets, almost every bid is accepted by the region and there are only a few extra households that get support in the later years. What is key is that at the mid-level budgets, in the early time steps of the simulation, some good, but rejected, households receive only national incentives, and cannot receive regional ones as well at a later time. Therefore, there is more of the regional budget left to fund later applications that may have been “outcompeted” earlier.

For the first come, first serve this dynamic does not occur because the regional government gives away funding as quickly as possible and there is less scope to receive only national incentives and free up regional budget for later time steps.

Simply put, funding everybody from the regional budget is not necessarily a good idea. At least some of the applications for regional funding would install also with just national funding. The money saved in the regional budget can then be put towards later applications that would not install without regional support, resulting in an overall higher installed capacity.

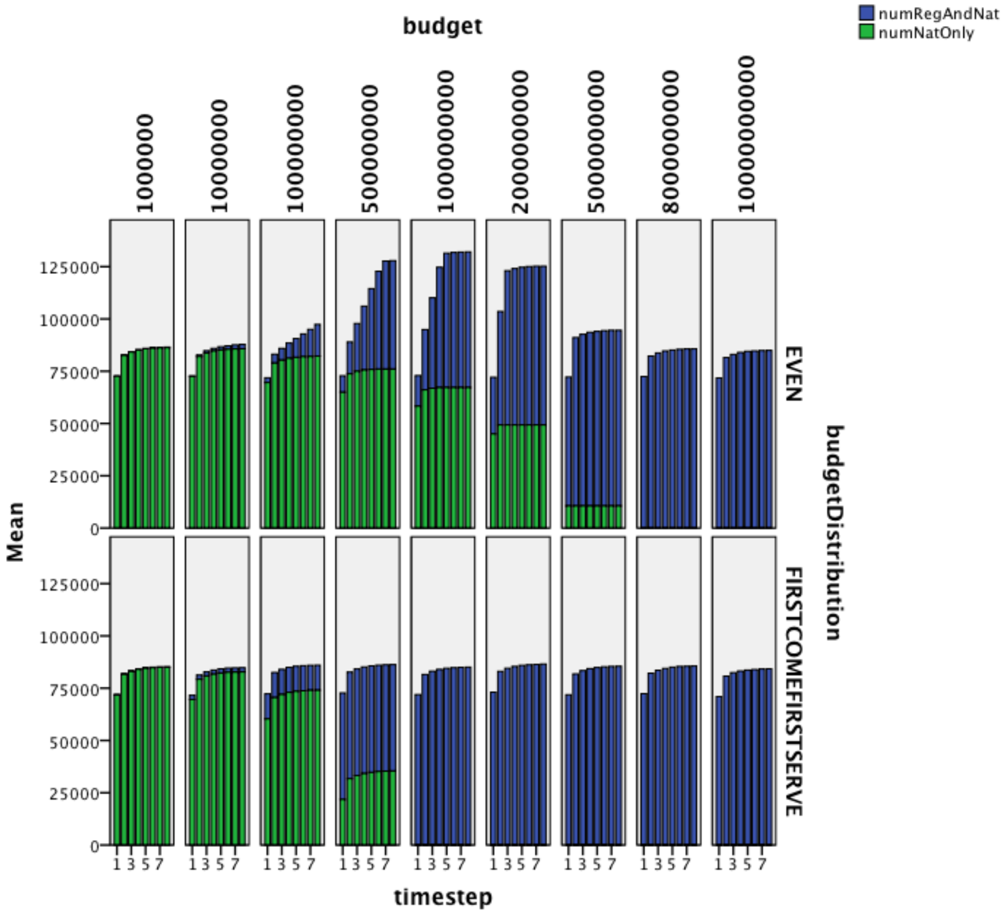

Figure 7: Mean number of installations over budget and time step by budget distribution.

Figure 7 shows empirical evidence that supports this intuition. It becomes obvious that that for first come, first serve budget allocation the number of participants does not decrease beyond a certain budget level, but that it does for even budget distribution. At the same time, the figure illustrates how the fraction of households that are funded by only national incentives (green bars) decreases as the budget for regional incentives is increased.

7 Optimising budget allocation over all time steps

The current design of the incentive design component optimises the allocation of funds to bid for support for each time step of the simulation individually. That is, for each list of bids received at the end of a time step, the optimal allocation of funds is computed without considering subsequent time steps. This approach is not necessarily optimal with respect to the overall outcome at the end of the simulation. In particular, the funds allocated to a particular bid in an early time step may have been better allocated to a similar bid with a higher power to cost ratio at a later time step. Indeed, we have observed this phenomenon when comparing the optimisation approach with an ordering heuristic [5].

The difficulty with making a globally optimal decision is that during all but the final time step, we have only partial information available. We do not know how many bids will be made in subsequent time steps, or what installation size they will offer for what requested funding. Instead of optimising known quantities, we have to estimate and predict future behaviour and make decisions based on those uncertain predictions.

Below, we describe our preliminary approach that allocates only a part of the total budget available at a time step to bids.

7.1 Approach

We assume that we will receive a set of bids, where each bid is of the form $B_j = (w_j, v_j)$, and where w_j is the cost, or grant requested, and v_j is the power which the bidder is offering to produce. We have a given total budget, and the objective is to choose a subset of the bids with maximal total power, subject to the constraint that the total cost of chosen bids is less than the budget. Thus, if C is the set of bids chosen (accepted), then we are maximising $\sum_{B_j \in C} v_j$ subject to the constraint that $\sum_{B_j \in C} w_j$ is no more than the budget. If we receive all the bids at the same time, this is an instance of the classic knapsack problem (cf. Deliverable 5.2 [2]).

The complication is that in the ePolicy optimisation problem we do not receive all the bids at the same time (nor do we receive bids one at a time, as in an online knapsack problem). We accept a collection of bids each year, over a number of years, and we choose the bids we will accept for each year.

A difficulty in solving this form of online problem is the division of the allocation of budget between the years. For instance, if we place no limit to the budget we use in the first year, then we may well use the whole budget then, which will prevent us from accepting any very attractive bids in later years. We can instead divide the budget allocation equally between the years, but this can also be sub-optimal, if, for example, we receive many more very attractive bids in the first year than expected. The idea behind our approach here is to take into account any information we might have about the future bids we expect to receive, and use this to determine whether a current bid is worth accepting, or if we expect to be able to do better later.

7.2 Basic Algorithm

Here we describe the basic structure of our algorithm, as applied to the bids received in a particular year. The idea is to try to choose the bids that have maximal ratio of power to cost, taking into account the expected future bids.

We will use a function V that, given remaining budget R, estimates the total value (total power) $V(R)$ we expect to achieve in later years. A single year's bids $\mathcal{B}$ consists of a set of pairs of the form $B_j = (w_j, v_j)$, where w_j is cost to the budget, and v_j is the value (power gained). We first sort $\mathcal{B}$ in such a way that it has decreasing values of $\frac{v_j}{w_j}$, breaking ties by increasing value of w_j, i.e., for $i, j \in \{1, \ldots |\mathcal{B}|\}$, if $i \leq j$ then $\frac{v_i}{w_i} \geq \frac{v_j}{w_j}$, and if $\frac{v_i}{w_i} = \frac{v_j}{w_j}$ then $w_i \leq w_j$.

Suppose at any point in the algorithm (applied to that year) we are considering bid $B_j = (w_j, v_j)$ and suppose we have remaining budget E. If we accept bid B_j we expend cost w_j to gain value (power) v_j.

Suppose we stop without accepting bid B_j. Then the expected extra total value we will obtain using the future years' bids is $V(E)$. If, on the other hand, we accept bid B_j and stop then, the expected extra total value we will obtain using the future years' bids will be $V(E - w_j)$. Thus it is better to (a) accept bid B_j and then stop, rather than (b) stopping now, rejecting bid B_j, if and only if $v_j + V(E - w_j) > V(E)$, i.e., $v_j > V(E) - V(E - w_j)$. This is the idea behind the test in the algorithm which determines whether to accept bid B_j.

Let E be the current remaining budget as we receive the current year's bids. At each point, R will represent the remaining budget.

Basic Algorithm

$R := E$

Total_Value := 0

for $j = 1, \ldots, |\mathcal{B}|$,

    if $v_j > V(R) - V(R - w_j)$ **then**

        (* We accept bid $B_j = (w_j, v_j)$ *)

        Accepted_Bids := Accepted_Bids $\cup \{B_j\}$

        Total_Value := Total_Value $+ v_j$.

        $R := R - w_j$

        end (* **if** *)

    end (* **for** *)

Naturally, the final value of variable Accepted_Bids will contain the collection of accepted bids, with total power equalling (the final value of) Total_Value, which equals the sum of v_j over all j such that B_j is in Accepted_Bids.

Roughly speaking, we are accepting bids whose ratio $\frac{v_j}{w_j}$ is greater than we would expect to obtain from the worst bids we would expect to be accepting from future years.

7.3 Estimating value $V(R)$

The basic algorithm involves the key test of whether $v_j > V(R) - V(R - w_j)$. In this section we discuss our approach for estimating $V(R)$ (and hence also $V(R) - V(R - w_j)$). Recall that $V(R)$ is the expected additional power we can achieve using the bids in future years, given the remaining budget R.

An important consideration is the efficiency of determining if $v_j > V(R) - V(R - w_j)$. The size of the collection of bids $\mathcal{B}$ will often be very substantial, and the basic algorithm loops over every element of $\mathcal{B}$. We thus may need that the test $v_j > V(R) - V(R - w_j)$ is not too expensive.

7.3.1 Estimating $V(R)$ based on collection $\mathcal{S}$ of future bids

First let's consider how we could proceed if we knew the collection $\mathcal{S}$ of bids we will receive in the future. Let $V^{\mathcal{S}}(R)$ be our estimate of $V(R)$ for this situation. For reasons of efficiency, we simplify in a couple of ways. We define $V^{\mathcal{S}}(R)$ to be equal to the maximum total power achievable given set of bids $\mathcal{S}$ and total budget R, where fractions of bids are allowed. If a fraction p between 0 and 1 of bid (w_j, v_j) is accepted then it costs pw_j to the budget, and gives us an additional pv_j of power. Thus firstly, we are ignoring the fact that the future bids $\mathcal{S}$ will be split between years; and we only know which bids we receive one year at a time; and secondly, we approximate the situation by allowing fractional bids. The second assumption will not likely make much difference to the result, apart from in some exceptional circumstances (and the fact that $V^{\mathcal{S}}(R)$ will then depend more smoothly on R than if we were not to allow fractional bids, may in fact help the accuracy of the overall algorithm). The first assumption will tend to mean we are somewhat overestimating $V(R)$ by $V^{\mathcal{S}}(R)$.

A straight-forward implementation of the $V(R)$ within the algorithm will still be rather expensive, with each call of $V^{\mathcal{S}}(R)$ being linear in $|\mathcal{S}|$. However, we have implemented a more sophisticated version of the algorithm which avoids this, by computing $V^{\mathcal{S}}(R)$ (and $V^{\mathcal{S}}(R - w_j)$) incrementally, so that typically only a small number of operations are needed for each loop index j.

7.3.2 Estimating $V(R)$ based on a distribution over future bids

Now, we consider the case where we have an estimate Q of the distribution governing future bids, and that we expect N future bids for some natural number N (which may well tend to decrease over the years).

We use a Monte-Carlo simulation algorithm involving a number M of trials. For each $i = 1, \ldots, M$, we create a random collection of bids $\mathcal{S}_i$ of cardinality N, drawn independently with distribution Q.

We then estimate $V(R)$ as $\frac{1}{M} \sum_{i=1}^{M} V^{\mathcal{S}_i}(R)$, where $V^{\mathcal{S}_i}(R)$ is defined below.

Generating the distribution Q

A simple approach is to let Q be the distribution of bids we have seen so far. So, for the first year, Q is just the collection of bids submitted in the first year. For the second year, Q is the collection of bids submitted in the first two years.

If we have additional information about what Q might be like, we could add extra elements to take this into account, or use some weighted average between the bids we have so far, and some guessed distribution.

If the distribution changes over time, so that e.g., the bids for the second year are much less attractive than those for the first year, then this should be taken into account. We might, for example, give more weight to the most recent bids.

More details on the approach can be found in [8].

7.4 Preliminary experimental results

We ran a set of preliminary experiments to compare this approach to the pure optimisation approach. We used the budgets levels detailed in Section 4, with both national and regional incentives enabled and first come, first serve budget distribution for the Emilia-Romagna region.

Figure 8 compares the total installed capacity at the end of each simulation. The results achieved by both approaches are very close, with the partial allocation approach achieving higher installed capacity in the majority of cases. There are cases where lower installed capacity is achieved as well. These results demonstrate the uncertainty inherent in this approach – we do not know future bids and can only make assumptions what they may be. Such assumptions may turn out to be wrong.

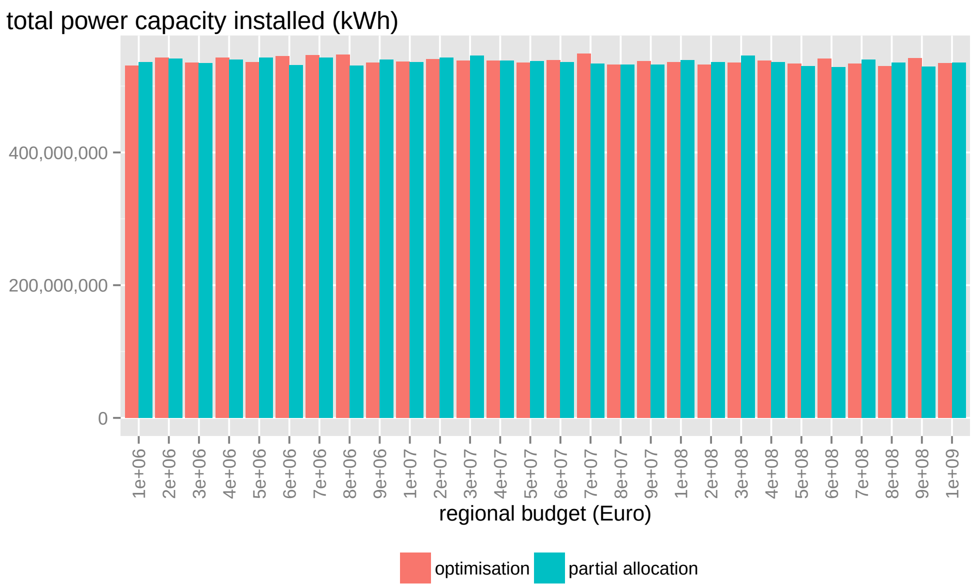

Figure 8: Total power capacity over regional budget for optimisation and partial allocation.

Nevertheless the results show the promise of the approach. We are able to improve on results that are provably optimal for each single time step by taking the entire simulation into account. This is very desirable, as it enables the policy maker to achieve more with the same budget.

Figure 9 presents the same comparison for the total cost incurred. Again the total cost is very similar for both approaches (except for the three largest budgets) and there are cases where the partial allocation approach spends less than the optimisation approach, but achieves a higher installed capacity. This again underlines the promise of the approach – we are able to exploit cases where not funding a bid not but a better one later will yield an overall better result.

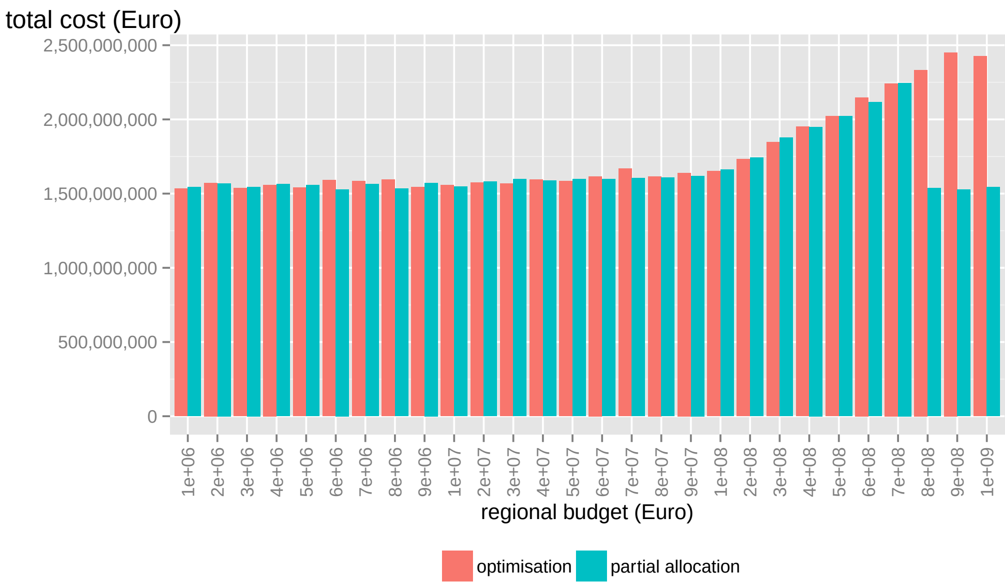

Figure 9: Total cost over regional budget for optimisation and partial allocation.

For the largest three budgets, the overall cost of the partial allocation approach is much lower than that of the optimisation approach, even though the achieved installed capacities are similar. Indeed, in two out of three cases, the installed capacity achieved by the partial allocation approach is slightly higher.

It is unclear why this is happening and what dynamics contribute to this phenomenon. Further research is needed and more work required to make it feasible to include this approach in the ePolicy system. The aim of the preliminary implementation and evaluation here is to highlight the promise of the approach as a potential direction for future research, not to present an alternative system ready for deployment.

8 Conclusion

We have demonstrated the capabilities of the integrated social simulator and incentive design component. We have explored a wide range of different values for the scenario parameters. In particular, our evaluation focused on the effectiveness of the different policy instruments a policy maker may employ to incentivize individuals to install solar power.

In addition to investigating the behaviour of the incentive design and simulator component on the Emilia-Romagna region, we ran the same experiments for the Bologna region. The results for both regions are qualitatively very similar, with the overall magnitude smaller for Bologna. This is to be expected, as Bologna is only a part of the Emilia-Romagna region. The fact that the results are qualitatively similar is encouraging, as it provides additional validation for the component.

Our experimental results demonstrate that regional incentives play only a minor role compared to the national incentives, although they seem to act as a catalyst. As the budget for regional incentives increases, the total installed capacity does as well if national incentives are present.

The most efficient incentive mechanism in terms of the ratio of cost to added installed capacity appears to be tax benefits. While it does not quite achieve the level of added capacity feed-in tariffs do, the cost it incurs is much smaller.

The aim of the experiments and their description here is not to derive predictions or recommendations for real world outcomes, but rather to show how a policy maker can explore the parameters relevant to her particular scenario and analyse the results. In addition, our evaluation can be used to identify the dynamics that affect the outcome of policy decisions.

We have furthermore offered some intuition, based on empirical evidence, to explain the effect that an increase in regional budget results in a decrease in installed capacity, as observed in deliverable 5.4. We also presented a preliminary approach and evaluation of a way of optimising the allocation of funds to bids over the entire simulation rather than for each time step separately. Such an allocation scheme is likely to avoid this decrease in installed capacity with increasing budget. The preliminary results are very encouraging.

References

[1] Tina Balke and Nigel Gilbert. The first complete version of the agent-based model (ePolicy Deliverable 4.1). Technical report, University of Surrey, 2013.

[2] Simon de Givry, Lars Kotthoff, Helmut Simonis, and Barry O'Sullivan. Prototype of incentive policy mechanism (ePolicy Deliverable 5.2). Technical report, Cork Constraint Computation Centre, 2013.

[3] Alan Holland and Barry O'Sullivan. Survey of game theoretic tools in dynamic environments for policy management (ePolicy Deliverable 5.1). Technical report, Cork Constraint Computation Centre, 2012.

[4] Peter George Johnson, Tina Balke, and Nigel Gilbert. Report on the policy instruments considered and their likely effectiveness (ePolicy Deliverable 4.3). Technical report, University of Surrey, 2014.

[5] Peter George Johnson, Tina Balke, and Lars Kotthoff. Integrating optimisation and agent-based modelling. In *28th European Conference on Modelling & Simulation*, Brescia, Italy, May 2014.

[6] Lars Kotthoff, Yulia Malitskaia, Barry O'Sullivan, Helmut Simonis, and Nic Wilson. Second prototype of incentive policy component (ePolicy Deliverable 5.3). Technical report, Cork Constraint Computation Centre, 2013.

[7] Lars Kotthoff, Yulia Malitskaia, Barry O'Sullivan, Helmut Simonis, and Nic Wilson. First demonstration of the testbed (ePolicy Deliverable 5.4). Technical report, Cork Constraint Computation Centre, 2014.

[8] Nic Wilson and Lars Kotthoff. Taking into account expected future bids in epolicy optimisation problem. Technical report, INSIGHT Centre for Data Analytics, 2014. Forthcoming.