ePolicy

Engineering the Policy-making Life Cycle

Seventh Framework Programme – Grant Agreement 288147

Deliverable 5.4
First demonstration of the testbed

Document type:	Deliverable
Dissemination Level:	Public
Editor:	Lars Kotthoff
Document Version:	1.0
Contributing Partners:	UCC
Contributing WPs:	WP5
Estimated P/M (if applicable):	3
Date of Completion:	31 Mar 2014
Date of Delivery to EC	31 Mar 2014
Number of pages:	14

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 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.

The ePolicy Consortium consists of the following partners: University of Bologna; University College Cork, National University of Ireland, Cork; University of Surrey; INESC Porto, Instituto de Engenharia de Sistemas e Computadores do Porto, Fraunhofer – Gesellschaft zur Foerderung der Angewandten Forschung E.V.; Regione Emila-Romagna; ASTER – Società Consortile per Azioni; Università degli Studi di Ferrara.

Authors of this document:

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

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

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. 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, and specifically the effects of the regional incentive budget (in Euro), the budget distribution (i.e. even, ramp-up, ramp-down, first come first served) and the regional objectives (i.e. max participation, max power, min cost), on the power produced in the region, the costs to regional government and the number of households installing PV.

The main change of the integrated social simulator and incentive design component compared to deliverable 5.3 is the inclusion of an updated social simulator. Additional verification has been performed on the simulator and the integrated model has also been through further validation, primarily using the data collected and presented in deliverable 2.1 on the regional energy plan. Furthermore, the simulator now comprises a stochastic element. That is, different runs of the same simulation scenario do not necessarily produce the same results.

It is for this reason that the individual simulations need to be repeated several times to allow to gain confidence in the results and estimate the uncertainty. In this deliverable, we explore the effect the number of repetitions has and conclude that it is sufficient to repeat each simulation run five times.

We present a set of comprehensive experiments that explore a range of different parameters and a total of 1200 different scenarios. In particular, we investigate a wide range of power targets and regional budgets. Furthermore, different budget distributions, optimisation objectives, and combinations of incentives are considered.

The results of our experiments demonstrate that the different funding scenarios have a significant impact on the outcome of the integrated simulation and optimisation runs. In particular, the installed power capacity can be improved by up to about 50% with the right incentives and budget.

Our experiments show that increasing the budget available for incentives helps to increase the installed capacity only up to a certain point. It is not the case that adding more funds increases outputs indefinitely. The most effective policy instrument for maximising power appears to be grants, with interest rate support marginally less effective. Using both incentive instruments at the same time does not seem to have a positive effect.

This deliverable and deliverable 4.3 share some of their contents because of the integrated nature of the social simulator and incentive design component (described in both deliverables and D5.3). In particular, the experimental setup and evaluation of results is shared in these deliverables. It is important to note that as the two components were originally developed separately, they are still treated as separate in evaluation and contingency planning as set out in deliverable 9.2.

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, 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 bids the agents generate are collected at the end of each time step. 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 objectives, 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. Following that, we investigate the specific integration issue arising from the update of the social simulator to include a stochastic component in Section 5. 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 6, we conclude in Section 7.

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 will be published in [5].

The main changes compared to the previous version of the integrated social simulator and incentive design component (deliverable 5.3 [6]) are as follows.

- The simulation includes a stochastic element that may give differing results for different runs of the same simulation.
- Additional verification and validation has been performed.

The changes in the social simulator part are described in detail in deliverable 4.3 [4].

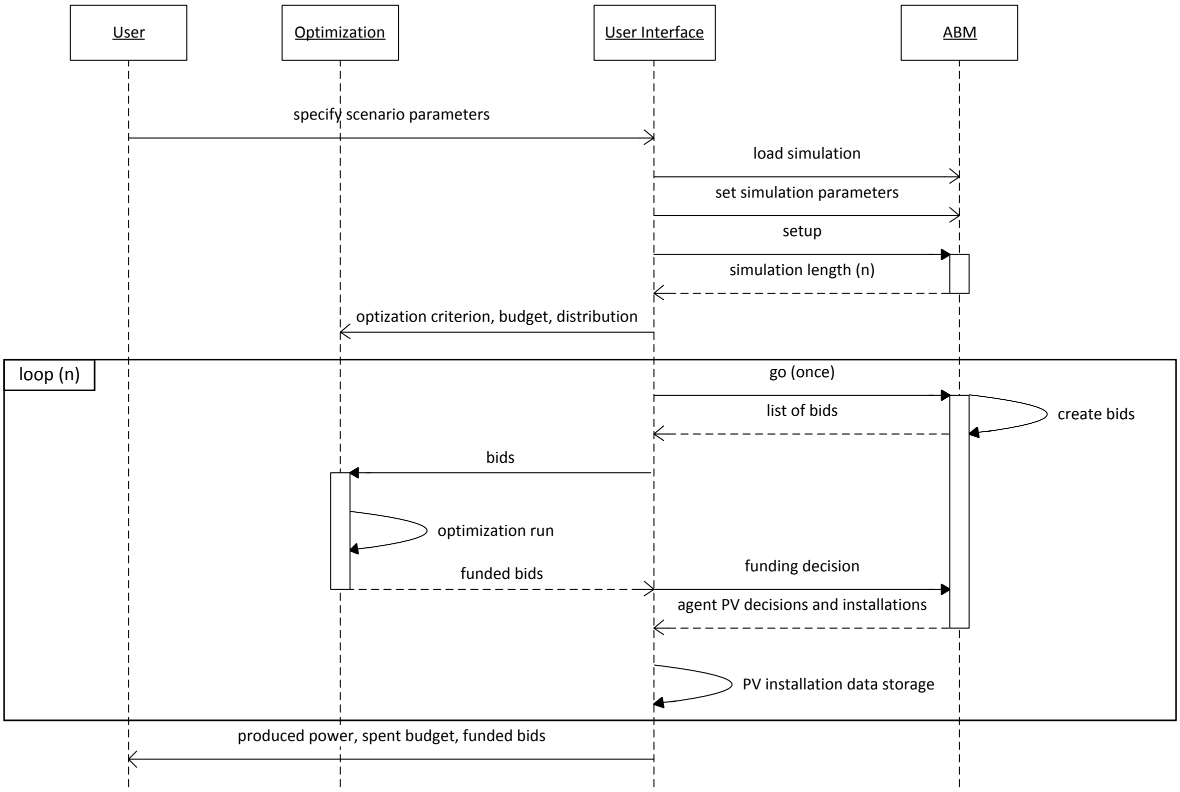

Figure 1: Sequence diagram of the integration of the social simulator and incentive design [5].

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 1200 different scenarios for the demonstration. The simulation for each scenario was repeated five times, as determined above. The scenarios evaluate the effects of the interactions of the key input parameters and the various levels they can be set at. In detail, the parameters we considered were the following.

Grants on and off

Interest rate support on and off
Budget Distribution first come first served, ramp-up, ramp-down, and even
Objective maximise power, maximise participation, minimise cost
Regional Budget €1,000,000, €10,000,000, €100,000,000, €1,000,000,000, and €10,000,000,000
Target Power 11,628,000,000 kWh, 23,256,000,000 kWh, 34,884,000,000 kWh, 46,512,000,000 kWh, 58,140,000,000 kWh, which correspond on average to 1,000 ktoe, 2,000 ktoe, 3,000 ktoe, 4,000 ktoe, and 5,000 ktoe

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.

5 Determining the number of repetitions

To determine the number of times we need to repeat each scenario to get reliable results, we ran the scenarios with the following four sets of parameters.

Grants on
Interest rate support on
Budget Distribution first come first served
Objective maximise power
Regional Budget €1,000,000 and €10,000,000,000
Target Power 11,628,000,000 kWh and 58,140,000,000 kWh

Each of these four scenarios was run 50 times. From this distribution of 50 samples per scenario, we sampled 100 different subsets of sizes 3, 5, 10 and 20 at random without replacement. We used the non-parametric Wilcoxon statistical test to compare the distribution of the original 50 values to the distributions of the smaller sample size. For each comparison, we recorded the probability that the two series of numbers were drawn from the same distribution.

The intuition behind these experiments is that if the difference between a smaller and a larger sample is statistically insignificant, it is statistically safe to use the smaller sample instead of the larger one. This reduces the number of repetitions for each scenario and the overall computational cost.

The difference between two distributions is said to be statistically significant if the probability that the same distribution is underlying both, as reported by a test such as the Wilcoxon test, is less than 5%. Table 1 shows the number of cases in which this was true for the different sample sizes and the key output variables.

sample size	3	5	10	20
total power capacity installed	4.5	3.5	1.25	0.5
total cost	4.25	2	1.25	0.25
total number of installations	4.25	2	1.25	0.25

Table 1: Probability (%) that the original distribution with 50 repetitions and the sampled distribution are statistically significantly different across all scenarios.

For 20 repetitions of a scenario, the probability that the difference between the measured outcomes and those measured for 50 repetitions is not statistically significant is very high (>99.5%). The probability decreases with 10 repetitions, with a 1.25% chance that the results are statistically significantly different (the opposite event). At 5 repetitions, the average probability of statistical significant difference increases to about 2.5% and for 3 repetitions even further.

There is a tradeoff between the reliability of the results and the computational cost involved in obtaining them. Ideally, we would run a large number of repetitions for each scenario, but this is infeasible in practice because of the high cost. Based on the results in Table 1, we decided to use five repetitions for the following experiments. The probability of a statistically significant difference to a much larger number of repetitions is very small and so is the computational cost. The probability does not improve much when doing twice as many repetitions. On the other hand, decreasing the number of repetitions only slightly increases the probability a lot.

6 Experimental results

We average the achieved results over the five repetitions for each scenario. In this document, we focus on a high-level analysis of the results that demonstrate that the integrated social simulator and incentive design component is working correctly. For a more detailed analysis of the effect of individual scenarios and parameters, we refer the reader to deliverable 4.3 [4].

We start by analysing the relationship between the total capacity installed as a result of the scheme and the total cost for the incentives. Figure 2 shows this relationship. Each point in the graph is an individual scenario; there are 1200 points in total. The shape of the point denotes what kind of regional incentives were used, its colour the optimisation objective.

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. However, the spread of the cost is relatively small compared to the absolute values – the range of the values measured in the experiments is less than 10% of the total amount. The range for the installed capacity is much larger, demonstrating the impact of the policy instruments.

There is a clear "baseline" at the bottom of the graph that shows the development when the incentives have no impact. All of the scenarios that use no incentives (circles) can be found there, as well as the scenarios with a small budget. If none of the incentives are used, the optimisation objective does not make a difference. The scenarios that have a larger budget and use at least one of the policy instruments achieve significantly higher installed capacities.

In general, the scenarios that minimise the cost (blue shapes) incur a smaller cost, as one would expect. Maximising the participation (i.e. the number of funded bids) incurs the highest costs. However, the smallest median cost is incurred by scenarios that maximise the power capacity. This can be also be seen in Figure 3.

The different "levels" that can be seen in Figure 2 correspond to the different budgets specified for the regional incentives. This parameter clearly has the largest overall impact on the

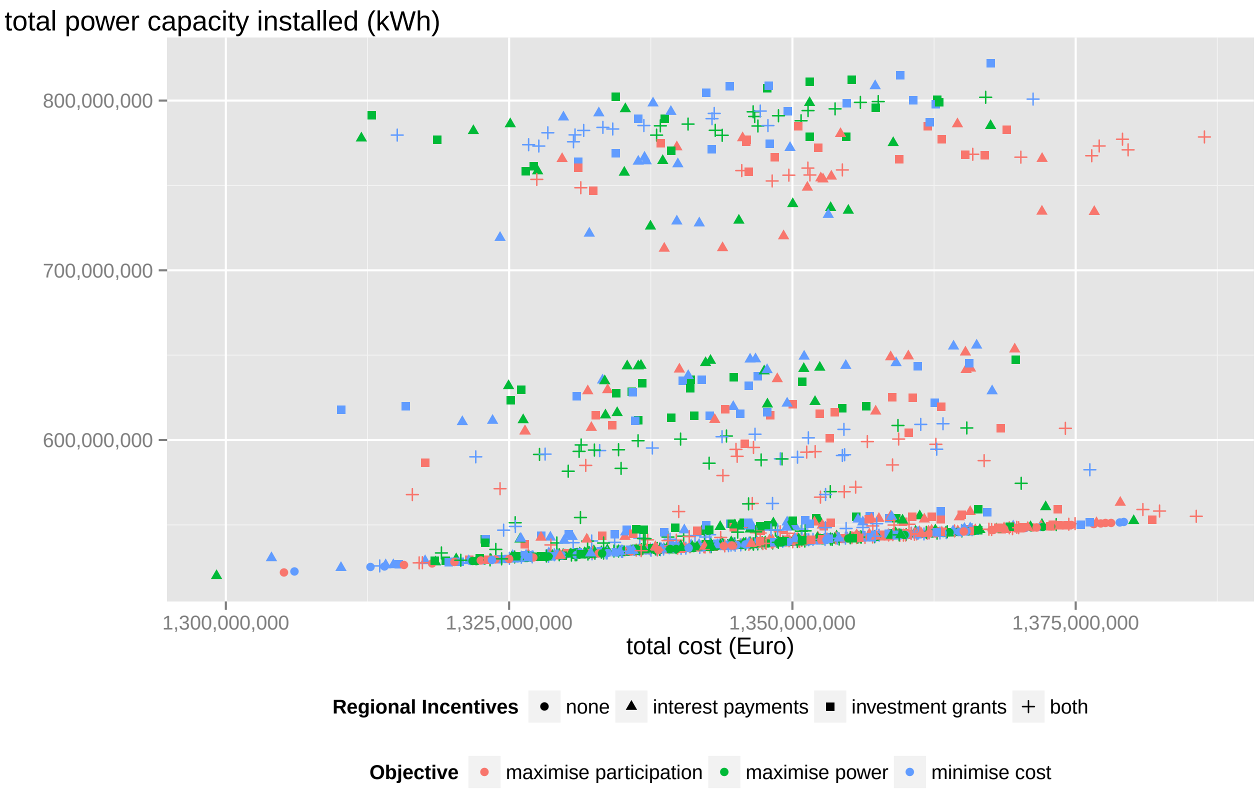

Figure 2: Total power capacity installed over total cost for all scenarios.

installed capacity.

There is no clear picture with respect to which incentive instrument has the larger impact. The scenarios that use investment grants appear to achieve a slightly larger installed capacity than the scenarios that use only interest rate payments, but the difference is small. However, using both at the same time does not seem to increase the installed capacity overall – it only increases the cost. Altogether, there is little difference in the total cost across the different combinations of incentives (Figure 4).

Figure 5 shows the relationship between the total number of installations after running a scenario to completion and the total cost. The graph has the same structure as Figure 2; each point represents a scenario, the shape denotes the kinds of regional incentives and the colour the optimisation objective.

Looking at the figure, it becomes immediately clear that the optimisation objective has a very significant impact on the number of installations. There is a clear separation between the scenarios that optimise for participation (red) and everything else. This is what we would expect intuitively – a lot of small installations would have the same power capacity as a single large one, but the large one would be more cost-efficient because of lower relative overheads such as installation cost. Asking the integrated component to maximise the number of installations (and thus the number of funded bids) biases the selection towards the smaller installations, of which many more can be funded for the same cost.

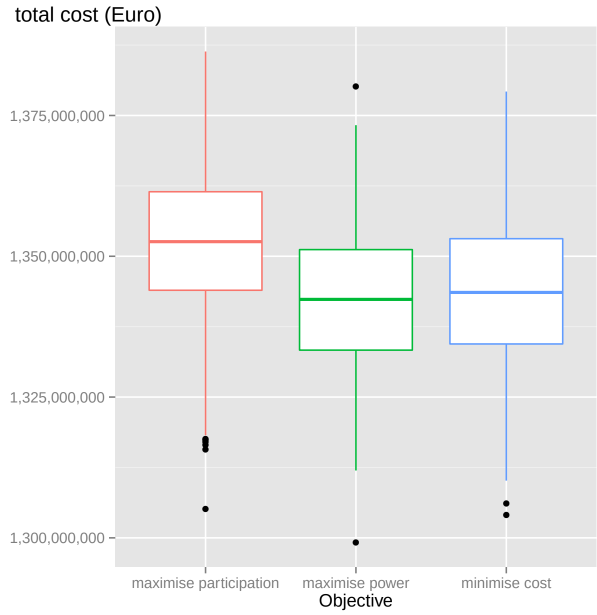

Figure 3: Total cost over optimisation objective for all scenarios.

For the same reason, the scenarios that optimise participation tend to be more expensive than the scenarios that optimise something else – achieving the same power capacity with several small installations is more costly than one large installation.

Note that the range of the number of installations the different scenarios cover is relatively small compared to their absolute number (about 10%). This underlines the importance of the national funding and other factors.

We now take a closer look at how the budget affects the installed capacity. Figure 6 shows this relationship. What becomes immediately clear is that increasing the budget only helps up to a certain point, after which it becomes detrimental. The highest installed capacities are achieved in scenarios which have a budget of €1 Billion. For most of the scenarios, the installed capacity is significantly higher than for any other scenarios with other budgets. The only exception are those scenarios that do not use any of the incentive instruments.

The difference between the two lowest budgets does not have a significant impact on the installed capacity, the scenarios for both budgets achieve more or less the same results. However, increasing the budget beyond €10 Million starts to show a noticeable impact. At that point, we also start to see a clear distinction between the scenarios that use incentive instruments and those that do not.

Increasing the budget to its largest value does not, as one may expect, increase the installed capacity even further, but in fact reduces it. It is unclear why this effect occurs and we will investigate this phenomenon in future work.

Figure 7 shows the relationship between the installed capacity and the incentives used directly and allows us to identify the most effective policy instrument. As mentioned above, using both interest payments and investment grants together achieves lower installed capacities than using just one of the instruments. It also makes clear that the incentives do have a

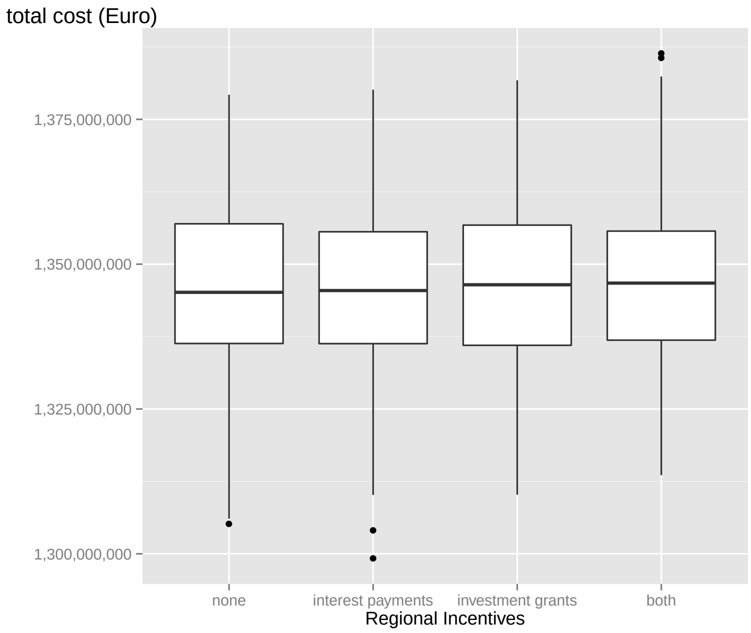

Figure 4: Total cost over regional incentives for all scenarios.

significant impact.

While the difference is small, using investment grants as incentive achieves better results than interest payments. Both in terms of the largest and smallest capacity installed across all applicable scenarios, giving investment grants is better. The difference in installed capacity with and without grant incentives is statistically significant as determined by the Wilcoxon test.

7 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.

One of the new challenges for this version of the integrated social simulator and incentive design component was that the simulation now contains a stochastic element that causes the results of different runs of the same scenario to be different. This means that in the integrated component, we need to run the simulation several times to get reliable results.

We established through experimentation that the best number of repetitions is five, presenting a good trade-off between the reliability of the results and their closeness to the actual distribution, and the computational overhead of repeatedly running the same scenario.

Our comprehensive experiments showed that regional incentives make a significant difference to the installed power capacity. With large budgets, the results are improved by as much as almost 50%. This result is very encouraging. Furthermore, we showed that the different parameters we have explored do have a significant impact on the outcome. This not only vali-

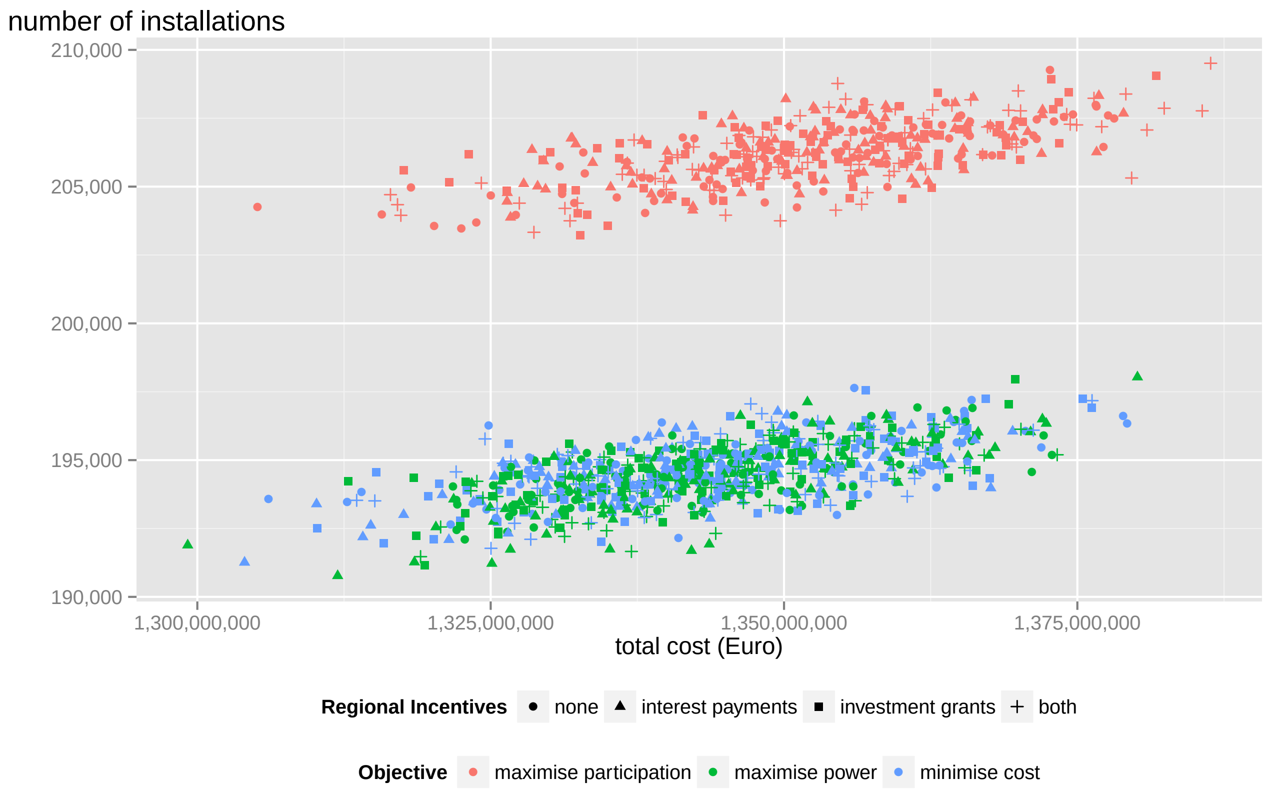

Figure 5: Total number of installations over total cost for all scenarios.

dates our choice of parameters, but also underlines the need for tools such as this component to help policy makers make their decisions.

The results show that at some point, increasing the budget available for incentives does not increase the overall power capacity installed. This make sense intuitively, as there is only a finite number of households who will consider installing PV. Once all of them have done so, additional funding is not going to have any effect. In our experiments, we have observed that the installed capacity actually decreases once the funding budget is beyond a certain point. We are planning to investigate the causes of this effect in future work.

The policy instrument that had the largest effect in our experiments was the grant incentive. Interest payments had a similar, but smaller effect. While using both incentive instruments together increased the cost significantly, it had no such effect on the installed capacity. This suggests that using only one of the two instruments investigated here may be preferable to using both.

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.

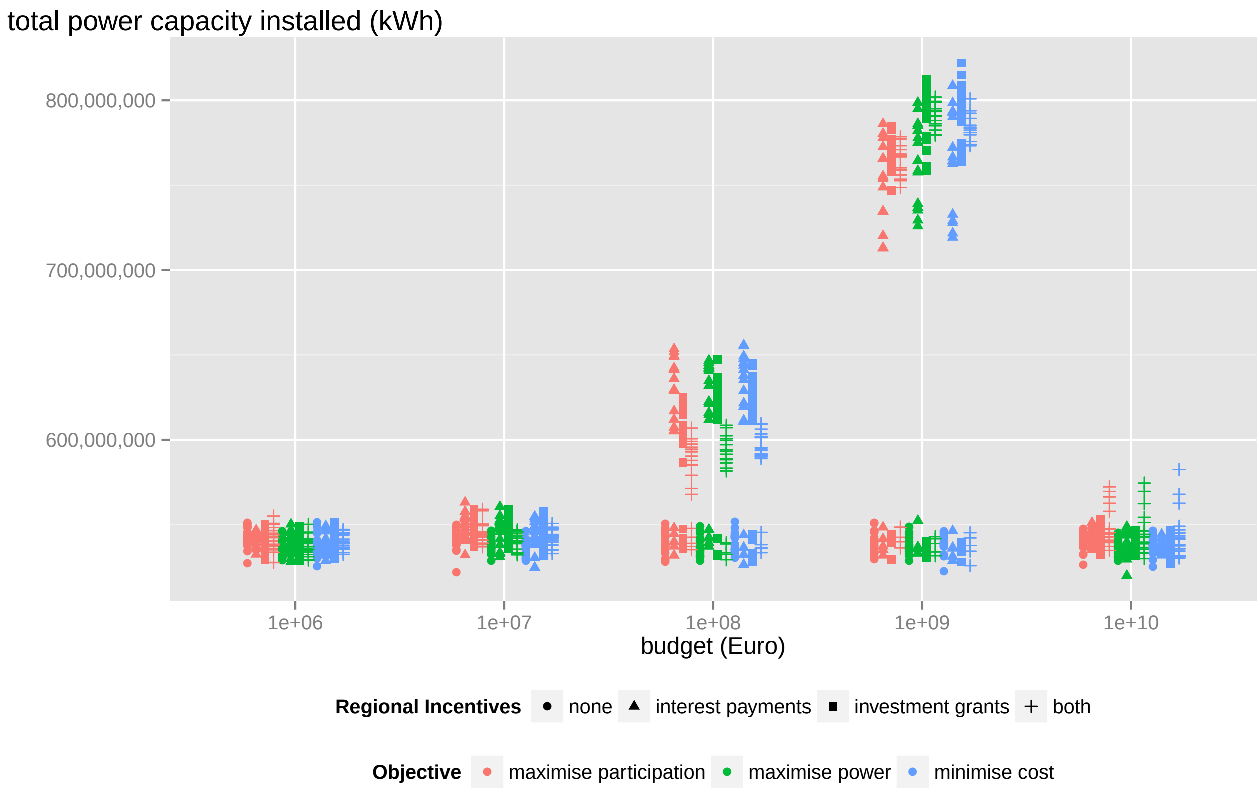

Figure 6: Installed capacity over incentive budget.

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.

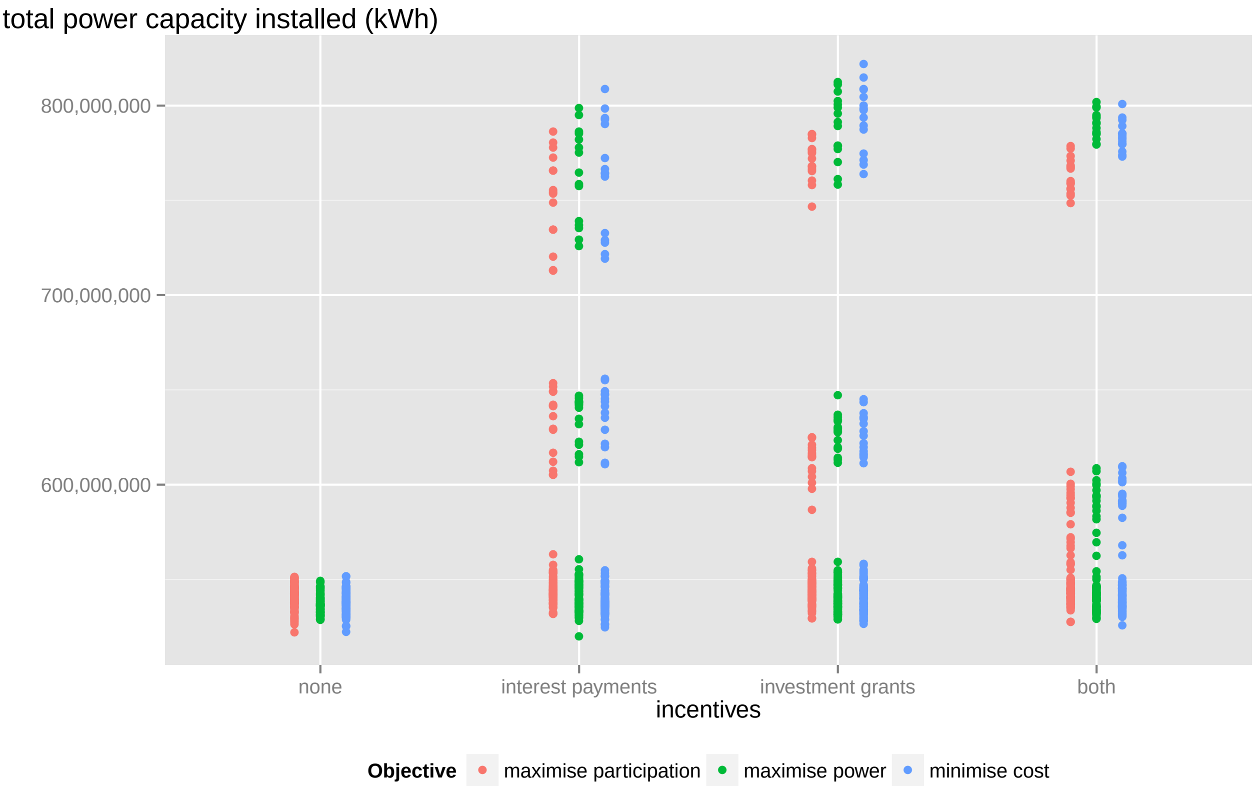

Figure 7: Installed capacity over incentives used.