Article:

Bell, A, Parkhurst, G, Droppelmann, K et al. (1 more author) (2016) Scaling up pro-environmental agricultural practice using agglomeration payments: Proof of concept from an agent-based model. Ecological Economics, 126. pp. 32-41. ISSN 0921-8009

https://doi.org/10.1016/j.ecolecon.2016.03.002

Scaling up pro-environmental agricultural practice using agglomeration payments: proof of concept from an agent-based model

Andrew Bell [a*], Gregory Parkhurst [b], Klaus Droppelmann [c], Tim G. Benton [d]

[a] Department of Environmental Studies, New York University, 285 Mercer St., New York, NY 10003, United States

[b] Department of Economics,Weber State University, 1337 Edvalson St., Ogden, Utah, 84408, USA

[c] PICOTEAM Ltd., Johannesburg, South Africa

[d] School of Biology, Faculty of Biological Sciences, University of Leeds, Leeds, United Kingdom

***corresponding author**

Ecological Economics 126 (2016) 32–41

Abstract

Rates of adoption of pro-environmental practices in agriculture in many parts of the world are low. In some cases, this is attributable to the private costs borne by farmers to adopt these practices, often well in advance of any benefits – public or private – that they may bring. Monetary incentives, such as through payments-for-ecosystem services (PES) programs, may be of assistance, and in this study we examine the potential for a recent innovation (the agglomeration payment) to improve adoption of pro-environmental practice in a rural agricultural context. Agglomeration payments include bonus payments for adoption by neighboring farms, which may help to encourage both compliance with the program they promote as well as the overall diffusion of the program across rural contexts. We develop an abstract agent-based model (ABM) of an agglomeration payment program to encourage adoption of the pro-environment practice of conservation agriculture (CA). We find that agglomeration payments have the potential to improve levels of adoption of pro-environmental practice per program dollar, and may help to reduce required spending on project monitoring and enforcement.

1 Introduction

In many parts of the world, agriculture leads to significant environmental impacts (Tilman et al., 2001). In some cases, mitigating impacts can also benefit farmers through gains in efficiency, such as the potential savings on water and other inputs from precision and conservation agriculture (Mondal et al., 2011; Palm et al., 2014). In other cases, reducing environmental impacts may lead to some significant costs. For example, in the case of improving biodiversity, land may need to be set aside, or managed in other ways, leading to drops in yields or income, heavier or different workloads, or opportunity costs in general. Other practices promoting environmentally beneficial agriculture, such as better soil management, may incur immediate costs, yet take time to develop private, as well as public benefits (Pannell et al., 2006). Given the social goods that may come from enhancing sustainability, yet the private costs of adopting, this presents something of a conundrum (Vanclay and Lawrence, 1994), and as a result, the uptake of pro-environmental management by farmers is often low.

To make environmentally-sound agriculture more attractive financial payments can and do play a role (e.g., Ndah et al., 2014). These may be a subsidy payment to compensate farmers for the costs, as typically articulated by the EU's Common Agricultural Policy (e.g., Donald, Pisano, Rayment, & Pain, 2002), or they may be articulated more positively as payment for ecosystem services (PES). The former (subsidy) framing may be targeted towards a broad bundle of behaviours; the latter (PES) framing, is perhaps more typically aimed at specific schemes, such as managing water quality, or, in the developing world, uptake of "conservation agriculture".

Despite three decades of developing PES schemes (Gómez-Baggethun et al., 2010) they remain characterized by low uptake rates, and where joining a scheme does occur, non-compliance with the expected management is often an issue. The former may arise from a variety of reasons, including incentives being insufficient, the payments coming from non-trusted sources, uncertainty about the length of the scheme and a variety of other social and cultural reasons (Engel et al., 2008; Ghazoul, 2007; Jack et al., 2008; Kroeger and Casey, 2007). Some of the same reasons affect compliance with the scheme and both lead to an escalation of scheme costs: either by requiring larger incentive payments or increases in the costs or monitoring or policing.

New incentive mechanisms are being designed that may address both problems simultaneously. From experimental economics, an innovative approach is to use agglomeration payments (Parkhurst et al., 2002) that offer payment to the adopter for different practices they adopt, but offer bonus payments if the adopter's neighbours also change their practices. Agglomeration payments have been explored in the literature largely as a spatial incentive and a means to obtain coordination in land use (Parkhurst and Shogren, 2007), with particular focus on biodiversity conservation, where contiguity in preserved habitat is of particular value (Albers et al., 2008; Drechsler et al., 2010, 1999). Agglomeration payments have the potential to exploit inherent informal institutions of social norms that regulate farmer's behavior within the farmer's community. Creating interdependencies between neighboring farmers' agricultural decisions has the potential to strengthen the community's organizational structure, promote information diffusion and transfer of technologies, and increase adoption and compliance of pro-environmental management. It may also aid compliance as the impact of a detected noncompliance not only results in the loss of the direct payment for the farmer in noncompliance, but also imposes a cost on all neighboring landowners engaged in the scheme who lose the agglomeration payment associated with the shared border. If informal institutions such as social norms increase the propensity for compliance (i.e., in order to avoid social disapproval (Balliet et al., 2011; Grasmick and Green, 1980) and the subsequent loss of social capital), then monitoring and enforcement costs will be smaller for an agglomeration bonus mechanism relative to other PES schemes such as a flat subsidy mechanism.

To our knowledge, there is only one evaluation of agglomeration payments in smallholder agriculture, and it is at an early stage of a short pilot to promote conservation agriculture in the Shire Basin, Malawi (Ward et al., 2015). With the long lag times from adoption to accrued benefits in pro-environmental behaviors like conservation agriculture, it is difficult to experiment empirically on different designs of PES schemes, whereas it is more straightforward to experiment *in silico*. To this end, we develop an abstract agent-based model (ABM) of a rural landscape of farmers making a choice about whether to adopt pro-environmental practices. ABMs are an ideal framework to model uptake of PES schemes as the unit of a decision-making agent can capture more of the complexity of decision making by individuals based on their context and local social interactions , as well as heterogeneity across farmers, than might be possible with an equation-based or representative-farmer model (e.g., Diederen et al., 2003; Gabriel et al., 2009; Langyintuo and Mekuria, 2008; Marra et al., 2003; Sutherland et al., 2012).

We model the choice for farmers as a decision between practices with different yield functions. Pro-environmental behaviours initially impose a yield cost, but over a period of several years this cost declines and turns into both a small yield benefit, and a reduction in sensitivity to climate. Our model is generic, with such a difference in yield functions between the non- and pro-environmental behaviours being expected in a range of situations (including integrated pest management, integrated soil management, organic farming, managing pollinator populations and so on). However, to ground this firmly in a familiar context, we focus on a concrete situation: conservation agriculture. This is a suite of practices (including minimal tillage, mulching of crop residues, and crop rotation or intercropping) (e.g., FAO, 2012; Ficarelli, Chuma, Ramaru, Murwira, & Hagmann, 2003;

Kassam & Friedrich, 2010; Marenya, Smith, & Nkonya, 2014; Ndah et al., 2014; Pretty et al., 2006), in which many of the private benefits of practice accrue only after some years of continuous adoption, a hard sell to small-holding farmers averse to uncertainty and high discount rates on the future (Pannell et al., 2014).

We model the decisions of individual farmers to adopt this practice as boundedly rational, risk-averse, maximization of expected utility, with farmers learning from others' experience and interacting to share potential benefits under an agglomeration payments program. We apply this model across a range of situations varying the details of the payment program, environment, and farmer conditions. Our key aim is to examine the potential advantages agglomeration payments may have in promoting adoption of environmentally sound agriculture, and in particular their impact on the overall cost-efficiency of the scheme. Our key finding is that agglomeration payments have the potential to improve levels of adoption of pro-environmental practice per program dollar, and may help to reduce required spending on project monitoring and enforcement.

2 Methods

Agent-based models (ABMs) have emerged in recent decades as a powerful tool to study complex systems where outcomes at the system level are strongly influenced by interactions among heterogeneous individuals (and/or heterogeneous contexts) within the system (Brown, 2006). This is particularly the case in many agricultural system processes, where decisions such as land-use changes (Deadman et al., 2004), land sells (Bell, 2011) or migration (Kniveton et al., 2011) are often strongly influenced by the subjective norms of those in one's neighborhood or community, and where yields may depend on specific details of the farm and its locality (German et al, in press). In an ABM, individuals in the system (such as farmers, in the agricultural context) are modeled as agents who take information from their environment and other agents in the system, and make decisions based on a set of rules programmed into the model to represent real-life decision processes. Landscape-level outcomes emerge from the concerted actions and interactions of individual agents.

For the current study, the ABM approach allows examination of the boundedly rational, multi-step decision-process of adoption across a heterogeneous landscape of expected-utility-maximizing farmers interacting with each other and responding to a spatial incentive (the agglomeration payment). This section outlines the structure of our model, the set of experiments undertaken using the model, and our subsequent analysis. The complete description of the model follows the ODD Protocol (Grimm et al., 2010, 2006), an emerging standard protocol for the description of agent-based models, and is included as Appendix A. The ODD Protocol includes diagrams for all process flow and scheduling in the model, as well as details of all sub-models included in the overall simulation. In the current section we provide a summary description sufficient to describe model assumptions and structure, and frame modeling results appropriately.

2.1 Agent-based model structure - Summary

We model a landscape of farmers, managing a portfolio of plots, in which plots may take one of three different land-use states: i) conventional practice, ii) conservation practice, or iii) 'cheating' – claiming to adopt conservation practice while remaining with conventional practice. The model captures two distinct time processes – a decision time step, and an interaction timestep – in order to capture variation in the extent to which farmers share information and consider their options. The decision time step is nominally one year, in which the state of the system is assessed (including yields) and farmers make decisions on how to manage their land in the following 'year'. The interaction time-step is nominally a sequence of within-year events where farmers interact with one another to share information.

2.1.1 Environment and Policy

Plots are subject to independent random variation in rainfall in time and space. All farmers are eligible to receive payments for the adoption of conservation practice, with payments taking two different forms whose levels are varied in different simulation runs. 'Base payments' are standard subsidies given to farmers proportional to the area under conservation practice. 'Agglomeration payments' are bonus payments given to farmers adopting conservation practice proportional to the number of neighboring farms also adopting conservation practice. Electing to 'cheat' on any plot brings the same program benefits as adopting conservation practice, but with a possibility of being randomly checked and having to pay a fine proportional to the 'cheating area' – i.e., area registered as conservation practice, that is actually under conventional practice. Conventional and conservation practices have different per-area yield functions, while cheating has the same yield function as conventional practice but with the chance of being assessed a penalty:

$$Y_{conventional} = (600 + 1 \cdot Rainfall)(1 + 1.2 \cdot Soil)(1 + 1 \cdot Efficiency) \qquad (1)$$

$$Y_{conservation} = \left(\frac{min(8, Seasons)}{8}\right)(700 + 0.7 \cdot Rainfall)(1 + 1 \cdot Soil)(1 + 1.2 \cdot Efficiency) \quad (2)$$

$$Y_{cheating} = (600 + 1 \cdot Rainfall)(1 + 1.2 \cdot Soil)(1 + 1 \cdot Efficiency) - (rand() < \ p_{catch}) \cdot Penalty \quad (3)$$

These yield functions are intended as simple and abstract; as modeled, yields for the encouraged conservation practice are slightly higher on average, less dependent on climate and soil quality, more dependent on farmer technical efficiency, and have a lag time to full yields (an example shown in Figure 1). Though a penalty is not strictly a component of yield, it is placed here for computational convenience so that uncertain factors that have some random variance (yields, penalties) are treated separately from factors that are certain (side payments, program payments). The policy environment is designed as a scheme of payments running for a finite time, and aiming to compensate the period when adopters are paying a direct yield cost. The expectation is that most adopters, if they adopt for the length of the scheme, will, at the schemes' end realize sufficient benefits to remain as pro-environmental practitioners.

We did not directly model monitoring costs; rather, we model monitoring via the two parameters of a) per-area penalty if caught 'cheating' and b) p_{catch}, the likelihood of a particular plot being monitored. For the purposes of qualitative comparison, a higher p_{catch} implies a higher investment in monitoring, and thus higher monitoring costs.

2.1.2 Interactions – Sharing information

Between each decision timestep there are n_{int} interaction timesteps. A farm participates in an interaction timestep (i.e., tries to talk to other farms and re-evaluates their preferred course of action) with probability $p_{participate}$. If the farm *i* participates, it then exchanges crop experience (history of net yields i.e. yield minus any penalty) of all of its plots with each other farm *j* in turn, with probability equal to the strength of the network link $p_{link\ i,j}$ between the two farms. This is to say, if two farms *i* and *j* interact during a given interaction time step, that interaction includes an exchange of all of their past experienced yields for the three land use options. Network link strengths $p_{link\ i,j}$ are randomly assigned at initialization and then rescaled such that network link strengths are stronger on average with closer neighboring farms than with those further away. After sharing information, the farm then re-evaluates its chosen land-use portfolio for the upcoming season (Figure 2).

Figure 2 indicates that for each other farm *j* that farm *i* has learned crop experience from, farm *i* constructs a similarity cue based on the variation in per-area crop yields when the two farms *i* and *j* took the same action (conventional, conservation, or cheat) anywhere within their portfolios in the same year. Farms with the lowest mean-square variation from farm *i* would have the greatest similarity, while those with the highest variation (as well as those with no shared experiences taking the same actions in the same years) would have the least similarity. This similarity is used as a weighting factor in the farmers decision (described in Section 2.1.4) – as farmers consider the possible utility from different portfolios and draw on their shared experiences to do so, they will weight estimates from more 'similar' farms much more heavily.

2.1.3 Interactions – Side payments

Farms that are currently practicing or plan to practice conservation practice or cheat in any of their parcels, and could benefit by agglomeration payments if other of their neighbors choose to adopt, may offer a side payment to their neighboring farms as an encouragement (where neighboring farms are those farms whose plots fall within a set radius of any plots of the identified farm). The side payment captures the possibility that farmers may invest in encouraging their neighbors to participate in the program if for no other reason than to increase their own returns (much as marketing tools such as groupon.com convey savings to consumers who can encourage others to participate with them). It is treated here as it has been elsewhere (e.g., M Drechsler et al., 2010; 2007) as a literal side payment of wealth between farmers, but in practice might include favors, shared risks, or social approval as mechanisms. Whether such side payments would manifest in a real agglomeration payment program is not known, so that our inquiry will distinguish the extent to which this mechanism drives our findings.

In the model, offers start as a random fraction of the value of one season's agglomeration payment, and are incremented by additional random fractions in subsequent seasons, until they reach a maximum of one season's agglomeration payment, or the neighbor chooses to adopt conservation practice or cheat. Standing offers from neighboring farms are considered by the farm during the estimation of expected utility from different land-use portfolios.

2.1.4 Farm-level decision

In each decision timestep, farmers implement a particular portfolio of land-use on their plots. Their choice of what to do in the next decision timestep is updated during each interaction timestep in which they participate, following a boundedly-rational, risk-averse, future-discounting expected utility model. Each farmer, as well as having a defined area, explicit in space to manage, has idiosyncratic (randomly assigned) attitudes to risk and discount rates.

The farmers' decisions are made boundedly rational by the simplification to a reduced set of alternative portfolios. Farmers do not choose for each plot individually, but rather, compare a random subset of possible portfolios (e.g., with 3 plots and 3 options there are 9 different possible portfolios), always including the particular portfolio (the default decision) they are currently doing in this comparison. For each of the portfolios the farmers select to evaluate, they draw n_{draws} random time paths of future yields (for n_{years} into the future), using in each time path the experience learned from their network, weighted by the similarity of the farms from which these experiences are drawn to themselves, to estimate yields for these portfolios. Put simply, farmers' estimates for future yields are based much more heavily on the experiences of other farms they deem to be more similar to them, than on farms they deem to be less similar (based on the similarity of historical yields to their own). Potential bonus payments from the program are added as appropriate to each of these time paths, as well as costs of converting from one land use to another, as appropriate. In the case of agglomeration payments, farms include offers that have been made to them of side payments. Additionally, farms estimate with probability, p_{adopt} , whether each of their non-adopting neighbors will enroll in the program in a given year, and account for the cost of nudging that adoption via side payment. These complete time paths are converted to a utility basis using the farmer's risk preference, and discounted to the present using the farmer's discount rate. The average across time paths is determined using the similarity metrics for the corresponding farmer to weight each time path. Farmers then choose to adopt the portfolio with the highest expected utility. Because the set of portfolios always includes the current portfolio (with transition costs, if any, already paid), it will remain the selected option as long as it remains the best option.

2.1.5 – Initialization

At initialization for all simulation runs in this study, a landscape of plots is randomly generated, with fraction f_{fill} of the landscape made up of plots with radii drawn from [$\mu_{r,plot}$, $\sigma_{r,plot}$]. Plots are randomly assigned into farms of [$\mu_{num\ plots}$, $\sigma_{num\ plots}$] each, and all other characteristics of the farm – risk behaviors, network link strengths, etc. (Table 1) are generated stochastically and assigned to the farms. This creates a landscape of farmers with heterogeneous decision parameters, working on farms that in turn differ in size and structure.

The model proceeds up to timestep $t_{pilot\ start}$ with all plots using conventional practice. From timestep $t_{pilot\ start}$ through $t_{pilot\ end}$, a subset of n_{pilot} farms are selected to 'participate' in a pilot project of the conservation practice. Participation of these farms is achieved by manipulating the perceived utility derived from conservation practice, such that they choose an option for their farm that includes either i) conservation practice or ii) 'cheating' on at least one of their plots; the perceived utility is held high for this option for the duration of the pilot so that these farmers observe yields from conservation practice and cheating and store them in memory. Through this mechanism, knowledge about the conservation practice as well as cheating is inoculated into the landscape.

2.2 Experiment Structure

The structure of our model captures mechanisms of diffusion via social contact. At the beginning of each simulation run, all farms are implementing conventional practice on all plots. Several decision timesteps into the simulation, a small pilot program is seeded in the model landscape – some number n_{pilot} farmers are granted sufficient additional encouragement (by inflating the expected utility of a portfolio including one or both of conservation practice or cheating) to maintain that portfolio for some $n_{years\ pilot}$ years. This 'pilot' seeds experience of the new options – adopting conservation practice, or cheating – into the landscape. Over time the experience from these pilot plots is shared with others in the landscape, and where it leads others to try conservation practice, the innovation spreads.

The base case in our experiment includes the initiation of this pilot program, but no further subsidy is provided. For comparison against this base case, our *in silico* experimental design systematically varies four key policy variables (level of base subsidy, level of agglomeration payment, penalty for being caught cheating, and likelihood of being caught cheating) in a full factorial sweep (Table 1).

Additionally, as there are a large number of environmental, cultural, and social variables in our agent-based model that are likely to vary widely, we overlaid this systematic sweep of the policy variables across random Monte Carlo variation in the environmental and social variables (an early sensitivity analysis included as Appendix B identified 26 variables of interest, summarized in Table 1). Specifically, for each unique Monte Carlo set of 26 variables drawn (a single value drawn for each of the 26 variables from uniform distributions, with ranges defined in Table 1), a full sweep over the 4 policy variables was performed. In total, limited only by time on a shared high-powered computing cluster, we undertook 624 Monte Carlo sets, each including a 6x6x4x3 sweep of the policy variables, for a total of 269,568 modeling runs.

Choosing parameter ranges for a Monte Carlo experiment in an abstract model is challenging. In our case, there are few variables that have meaningful bounds in the literature, nor would it necessarily appropriate to peg some values to literature bounds while other aspects of the model structure did not conform to the system to which that literature was referencing. We chose parameter ranges that meaningfully scaled against the three yield functions (equations 1-3) for the land use options, which are treated as fixed. For parameters that did not tie directly to the yield functions (such as the size and number of plots, or the risk behaviors of farms) we attempted to choose ranges that would jointly create landscapes with wide variation in the number of farms; the range from small to large farms; or the behavior from low to high risk aversion among farms. In the results and discussion that follows we refrain from making inferences specific to particular environmental conditions, and instead present findings that emerge as robust over the range of Monte Carlo parameters tested in the study.

2.3 Analysis

We generate response surfaces for the key outcome of interest (area under adoption) as functions of various sweep variables (base payments, agglomeration payments, likehood of being caught, and the per-area penalty for cheating) as well as other outcomes (overall program costs).

To evaluate the importance of particular variables in affecting outcomes within the MC exercise, we estimated variable importances using the RandomForestClassifier in

Python (Sci-kit Learn, 2015). The RandomForestClassifier (rF) is a machine learning algorithm that builds regression trees using random subsets of the data, and assesses the goodness of fit of the tree based on its ability to predict the subset of the data that were not used to build the tree. Importantly, comparison of the random set of trees allows the importance of different variables to be assessed: ones that radically affect the ability of the tree to predict correctly are more important. Unlike conventional regressions, in which interactions of interest must be identified as part of the model, interactions are implicit in any decision tree developed within a random forest (i.e., a variable used to classify data at a particular node is implicitly interacting with the variable used to classify data at the previous, higher node). Due to the memory constraints invoked by the large dataset, we performed a composite Random Forest analysis using 50 different samples of 25000 data points (~10% of our dataset per sample) and analyzing using a classifier with 1000 trees, all other settings default.

3 Results

In a typical simulation run, all land is under conventional practice at timestep 0 (Figure 3A), with an initial pilot of the conservation practice (plus a few cheaters) launched in the first few years (Figure 3B), leading to adoption of conservation practice as new of conservation practice performance spreads (Figure 3C) with some disadoption as the direct program benefits expire (Figure 3D). An additional outcome of interest with agglomeration payments is often the level of spatial coordination or contiguity in adoption that is achieved. We measure contiguity here as the fraction of planted area in the neighborhood of a plot adopting CA that is also adopting CA – a similar measure to the mean proximity metric (Wu and Murray, 2008), but not discriminating plots by distance within the neighborhood, and also retaining a scale from 0 (all adoption is dispersed across separate neighborhoods) to 1 (all neighboring plots of an adopting plot are also adopters). We see agglomeration payments having a similar impact on contiguity of CA as on net adoption of CA (note that both measures will converge at higher levels of adoption) (Figure A1, Appendix C, examining contiguity of CA adoption). This is consistent with previous results examining agglomeration payments for biodiversity conservation (e.g., Drechsler et al., 2010; Parkhurst and Shogren, 2008), but is not the outcome of primary interest in this study where we examine effects on the cost-effective encouragement and diffusion of CA adoption.

3.1 Explaining model outcomes

We highlight four key model outcomes in this section. First, as might be expected, greater spending on program payments (either by increasing base payment or agglomeration payment levels) leads to greater levels of adoption. Figure 4 shows contours of area under adoption as a function of both 1) overall program costs for payments (y-axis) and 2) the fraction of payments costs that go to agglomeration payments; the shape of this surface is irregular as payments spending is also a model outcome and not an input variable. Moving up the 'Payments Costs' axis (y-axis) in each of Figures 4A-C, the area under adoption increases. Importantly, the highest levels of adoption occurred in programs that included agglomeration payments. Second, some level of increased investment in monitoring effectiveness (compare contours in Figure 4A with Figures 4B or 4C) leads to greater levels of adoption at lower levels of spending on payments. These first two results provide face validity for the proper function of our model, but raise questions regarding cost effectiveness – how is money best spent across base payments, agglomeration payments, or effort in monitoring?

To answer this question, our third key result is that, when some monitoring effort is present, agglomeration payments can improve cost effectiveness. The contours in Figure 4 (most clearly in Figures 4B and 4C, with higher monitoring effort) indicate that equal levels of area under adoption are achieved with lower spending on payments costs, as the fraction of payments spent as agglomeration payments increases. For example, under medium monitoring effort, the cost of achieving a fraction of 0.29 of the landscape under adoption (Figure 4B, middle contour) with no agglomeration payments is about 0.3 (units), but this drops to 0.2 when half of all payments are agglomeration payments. Mechanistically, the reasons for this are two-fold. Most importantly, agglomeration bonuses (and to a minor extent side payments among adoptees) help to meet the reservation price of farms that value the conservation practices less, without a blanket increase in the subsidy. As participation by neighbors increases, the net value of adoption rises (through additional agglomeration bonuses as well as side payments offered by neighbors). This efficiency improvement from agglomeration payments was demonstrated by Drechsler et al. (2010) in the context of biodiversity conservation for butterfly habitat, and demonstrates the potential for agglomeration payments to help subsidy programs approach heterogeneous farm-tailored pricing. The other mechanism in place is that the peer effect of receiving a side payment is only a guarantee when the practice is actually adopted; a neighbor may or may not honor a side payment for a farm that cheats (whether they will or not is known to both farms). Thus, incentives to cheat are reduced relative to incentives to adopt when agglomeration payments are included in the program. The overall contribution of side payments is comparatively low in our results, accounting for at most around 4% of adoption in cases where program spending is low overall (Figures A2-3, Appendix C, comparing simulation results with and without the side payment mechanism); however, we note that side payments are capped at a maximum of one season's agglomeration bonus in our study, capturing uncertainty over how long a recipient might stay in the program. Thus, in contexts where farmers might believe their neighbors would stay enrolled longer (and thus be willing to offer more to encourage them), these mechanisms might have greater importance.

The fourth key result is that agglomeration payments have potential to substitute for monitoring efforts in such subsidy programs. Figure 5A shows contours of area under adoption as a function of 1) program spending on payments and 2) likelihood of being caught, generated from all modeled cases where only base payments were applied; Figure 5B shows the same for all modeled cases where only agglomeration payments were applied. The 'concave-up' or "L" shape of the equal-adopted-area contours in 5B (absent in 5A) indicates that equivalent levels of adoption can be achieved either by increasing the likelihood of catching cheaters (i.e., spending more on monitoring programs) or by spending more on agglomeration payments. This pattern is robust to programs that include base payments as well (Figures 5C,D), where again it is clear that these payments work synergistically to provide greater levels of adoption for the same overall program cost than either payment form alone (black dashed line across all panels). This result is preserved in the absence of side payments (Figure A4, Appendix C, comparing simulation results with and without the side payment mechanism) but is less clear and consistent. Whilst our abstract model does not estimate the costs associated with increasing the likelihood of catching cheaters, this figure demonstrates the proof of concept through which agglomeration payments may substitute for monitoring effort. In conditions where monitoring is expensive (as it may be for rural areas in developing countries), payment programs that include agglomeration payments may be able to improve program efficiency.

3.2 Key environmental variables

Our model attempts to model the complexities of decision making such as occur in real situations. It is therefore instructive to look at the range of variables and processes in order to understand whether there are key variables that drive decisions. To identify the relative importance of each of the variables in the model that contributes to the area under conservation practice and program cost (Figure 6) we used random forest to estimate variable importance.

To reduce computational overheads, the variables in the Monte Carlo analysis were selected using an earlier sensitivity analysis on model outcomes (included as Appendix B), excluding model parameters that did not initially appear to strongly shape outcomes. Nonetheless, Figure 6 indicates that there are many variables with similar importances, and this highlights the complexity of decision-making processes: adoption and program cost do not depend on a very small subset of variables but many variables including the farmer's characteristic attitudes to risk and discounting, social networks, farm characteristics and environmental conditions. As an example, the variable 'chanceAdopt' (capturing the perceived likelihood that one's neighbors might adopt in future, used in the estimating future utility from adopting and receiving payments) appears important to both adoption and program cost, highlighting the importance of designing social research instruments to be able to i) identify whether this perception does in fact shape this kind of decision, and ii) if so, measure it.

4 Discussion

An important goal of the sustainable intensification of global agriculture is to ensure that any intensification occurs alongside improvements in sustainability (Garnett et al., 2013), which requires the uptake of pro-environmental behavior. Farmer decision making is a highly complex process involving a range of criteria unique to the individual and his/her context (Edwards-Jones, 2006); partly because of this, the literature on how best to encourage uptake of pro-environmental behavior is itself complex, and clear messages about the "best ways" to encourage uptake are not apparent. To study this process in general, we developed an agent-based model to capture some realistic dimensions of a payment for ecosystem services (PES) scheme. In particular, we used this model to examine the efficacy of an innovative form of PES scheme, the agglomeration payment. Our key findings are that inclusion of agglomeration payments in a PES program i) increases absolute uptake, ii) can decrease overall payment costs, and iii) may have potential to reduce the demands of program monitoring. Our model incorporated a range of contextual processes and variables – our analysis revealed that a significant majority these were important, leading to a general finding that decision making (and hence, levels of adoption and program cost) are governed by a large number of context-specific variables. This indicates that there may not be simple leverage points on which to focus behavioral-change studies: information and its diffusion, farmer attitudes to risks, social networks, details of the farm and its environmental context are all important. A simple “knowledge deficit” model does not fit.

The key findings regarding agglomeration payments derive from the peer effect that the payments introduce – in large part through the shift in direct payments as neighboring farms adopt, and to a lesser extent through any kind of additional encouragement these adopting neighbors might see beneficial to provide. Drechsler’s (2010) work highlighted the efficiency gains possible via side payments, as well as the spatial coordination important for biodiversity conservation, but we demonstrate here that these peer effects have much more to offer to agricultural contexts. The cumulative agglomeration payments along with the neighbors’ encouragements themselves (the ‘offer’ made to neighboring farmers) act jointly as a diffusion mechanism to spread adoption beyond the set of adopters convinced by the base subsidy, contributing further to the higher uptake that agglomeration payments facilitate. Incentivising farmers to work cooperatively is something that is increasingly recognized as important in a range of situations where “upscaling” – the adoption of a large enough area to influence ecosystem services operating at large scales – is important to create the social benefit (e.g., Concepción *et al.*, 2008; DeFries and Rosenzweig, 2010; Hodgson *et al.*, 2010; Benton, 2012; Tscharntke *et al.*, 2012; Sayer *et al.*, 2013), and this provides a mechanism. Sometimes peer pressure works against uptake (Sutherland *et al.*, 2011) and agglomeration payments are a route to overcome this by targeting an individual’s incentives rather than rewarding communities. Upscaling can create a virtuous circle, not only in terms of delivering ecosystem services (like population or water management that work at landscape or catchment scales), but also by producing social benefits in terms of communities of practice that themselves can drive benefits. For example, in the UK, (Gabriel *et al.*, 2009) hypothesized that the spatial clustering of organic farms might initially be sparked by the environmental conditions, but when the clusters grow they may support the local development of markets, and further incentivize conversion to organic farming.

Furthermore, to the extent that i) side payments are significant relative to the direct agglomeration bonuses and ii) any farmers are reluctant to make side payments to neighbors for cheating (i.e., the encouragement to adopt is stronger than the encouragement to cheat), agglomeration payments can improve compliance with payment schemes, furthering the efficiency gains achieved via side payments. This effect may have greater relevance for smallholder agricultural contexts than for the large landholders who have up to now been the focus of much agglomeration payments work (e.g., Hartig & Drechsler, 2010; Parkhurst & Shogren, 2008; Parkhurst et al., 2002). Where the costs of monitoring and follow-up are expensive, the integration of agglomeration payments into a payments scheme can reduce the role of monitoring in maintaining compliance.

The ABM is a flexible model than can be easily adapted to a range of case studies. We included a range of variables that have some empirical support in determining (a) yields and (b) decisions. The variable importance analysis (using the random forest algorithm) indicated that contrary to our expectations that a few variables might be particularly important, the majority of the variables were. This reinforces the view that technological adoption is highly complex and so a full understanding, either in the general case or for a particular situation, needs to recognize this complexity (Edwards-Jones, 2007; Pannell et al., 2006) (Edwards-Jones, 2006)

Our model was notionally based on the adoption of conservation agriculture in developing countries, which is often hindered by the interdependencies between agricultural productivity on small farms, and social and cultural values, as well as access to resources. Overcoming these obstacles to CA adoption requires the strengthening of organizations beginning at the community level, to promote dissemination of information and the transfer of knowledge between farmers as they learn how to adapt CA techniques to the agricultural environment (Ficarelli et al., 2003). Further, social relationships have significant and complex effects on individual decisions within villages and communities in developing countries. Strong social relationships translate into social capital, which generates greater trust and reciprocity allowing individuals more access to economic resources and extended social networks within the individuals community. Ignorance to the impacts of social capital when implementing a CA (or any other similar) management plan can have negative impacts on adoption and compliance (Hoang et al., 2006). In developing countries, informal institutions that capture social norms regulating individual behavior within the community can be an effective avenue to managing ecosystem services (Jones et al., 2008). Our model demonstrated that agglomeration payments have the potential to act as a spark to strengthen diffusion and build social cohesion in this context.

5 Conclusions

We developed an abstract ABM to investigate generic decision-making on land-use practices by farmers, especially in the form of encouraging pro-environmental behavior via payment for ecosystem services. Specifically, we were interested in the role agglomeration payments could play in making PES schemes more effective.

We found that agglomeration payments were synergistic with conventional payment programs, leading to 1) greater levels of pro-environmental behavior per payment dollar spent, 2) greater overall levels of pro-environmental behavior, and 3) the potential for payment structure to reduce the need for monitoring efforts that can prove costly in remote smallholder contexts. These results were robust to a wide range of environmental and social conditions. Our model also reinforced that individual farmer decisions result from interplay among many variables – social, economic and environmental – and so it is difficult to highlight key leverage points for behavioral intervention.

The growing numerical support for agglomeration payments as a conservation tool warrants further exploration via pilot study. The numerical work presented here has informed the design of a pilot study for improved adoption of conservation agriculture in Malawi, currently underway.

6 Acknowledgements

This work was funded under two independent projects. The first of these 'Smart Subsidies for Catchment Conservation in Malawi' (NERC Reference NE/L001624/1) was funded with support from the Ecosystem Services for Poverty Alleviation (ESPA) programme. The ESPA programme is funded by the Department for International Development (DFID), the Economic and Social Research Council (ESRC) and the Natural Environment Research Council (NERC). The second grant was funded by the BASIS AMA Innovation Lab, which is supported by the US Agency for International Development (USAID).

7 Literature Cited

Albers, H.J., Ando, A.W., Batz, M., 2008. Patterns of multi-agent land conservation: Crowding in/out, agglomeration, and policy. Resour. Energy Econ. 30, 492–508.

Balliet, D., Mulder, L., Van Lange, P., 2011. Reward, punishment, and cooperation: a meta-analysis. Psychol. Bull. 137, 594.

Bell, A.R., 2011. Highly Optimized Tolerant (HOT) Farms in Rondônia : Productivity and Farm Size , and Implications for Environmental Licensing. Ecol. Soc. 16, 39.

Brown, D.G., 2006. Agent-based models, in: Geist, H. (Ed.), The Earth's Changing Land: An Encyclopedia of Land-Use and Land-Cover Change. Greenwood Publishing Group, Westport, CT, pp. 7–13.

Deadman, P., Robinson, D., Moran, E., Brondizio, E., 2004. Colonist household decisionmaking and land-use change in the Amazon Rainforest: an agent-based simulation. Environ. Plan. B Plan. Des. 31, 693–709. doi:10.1068/b3098

Diederen, P., Meijl, H. Van, Wolters, A., Bijak, K., 2003. Innovation Adoption in Agriculture : Innovators , Early Adopters and Laggards. Cah. d'economie Sociol. Rural. 67, 20–50.

Donald, P.F., Pisano, G., Rayment, M.D., Pain, D.J., 2002. The Common Agricultural Policy, EU enlargement and the conservation of Europe's farmland birds. Agric. Ecosyst. Environ. 89, 167–182. doi:10.1016/S0167-8809(01)00244-4

Drechsler, M., Johst, K., Shogren, J.F., 1999. An agglomeration payment for cost-effective biodiversity conservation in spatially structured landscapes. Ecol. Econ.

Drechsler, M., Wätzold, F., Johst, K., Bergmann, H., Settele, J., 2007. A model-based approach for designing cost-effective compensation payments for conservation of endangered species in real landscapes. Biol. Conserv. 140, 174–186. doi:10.1016/j.biocon.2007.08.013

Drechsler, M., Wätzold, F., Johst, K., Shogren, J., 2010. An agglomeration payment for cost-effective biodiversity conservation in spatially structured landscapes. Resour. Energy Econ. 32, 261–275.

Edwards-Jones, G., 2007. Modelling farmer decision-making: concepts, progress and challenges. Anim. Sci. 82, 783. doi:10.1017/ASC2006112

Engel, S., Pagiola, S., Wunder, S., 2008. Designing payments for environmental services in theory and practice: An overview of the issues. Ecol. Econ. 65, 663–674. doi:10.1016/j.ecolecon.2008.03.011

FAO, 2012. Economic aspects of Conservation Agriculture [WWW Document]. URL http://www.fao.org/ag/ca/5.html

Ficarelli, P., Chuma, E., Ramaru, J., Murwira, K., Hagmann, J., 2003. Strengthening local organizations for conservation agriculture - some experiences from South Africa and Zimbabwe, in: Proceedings of the 2nd World Congress on Conservation Agriculture. Iguaçu Falls, Paraná, Brazil, pp. 11–15.

Gabriel, D., Carver, S.J., Durham, H., Kunin, W.E., Palmer, R.C., Sait, S.M., Stagl, S., Benton, T.G., 2009. The spatial aggregation of organic farming in England and its underlying environmental correlates. J. Appl. Ecol. 46, 323–333. doi:10.1111/j.1365-2664.2009.01624.x

Garnett, T., Appleby, M.C., Balmford, A., Bateman, I.J., Benton, T.G., Bloomer, P., Burlingame, B., Dawkins, M., Dolan, L., Fraser, D., Herrero, M., Hoffmann, I., Smith, P., Thornton, P.K., Toulmin, C., Vermeulen, S.J., Godfray, H.C.J., 2013. Sustainable Intensification in Agriculture: Premises and Policies. Science (80-.). 341, 33–34.

Ghazoul, J., 2007. Challenges to the uptake of the ecosystem service rationale for conservation. Conserv. Biol. 21, 1651–2. doi:10.1111/j.1523-1739.2007.00758.x

Gómez-Baggethun, E., de Groot, R., Lomas, P.L., Montes, C., 2010. The history of ecosystem services in economic theory and practice: From early notions to markets and payment schemes. Ecol. Econ. 69, 1209–1218. doi:10.1016/j.ecolecon.2009.11.007

Grasmick, H., Green, D., 1980. Legal punishment, social disapproval and internalization of illegal behavior. J. Crim. Law Criminol. 71, 325–335.

Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., Goss-Custard, J., Grand, T., Heinz, S.K., Huse, G., Huth, A., Jepsen, J.U., Jørgensen, C., Mooij, W.M., Müller, B., Pe'er, G., Piou, C., Railsback, S.F., Robbins, A.M., Robbins, M.M., Rossmanith, E., Rüger, N., Strand, E., Souissi, S., Stillman, R. a., Vabø, R., Visser, U., DeAngelis, D.L., 2006. A standard protocol for describing individual-based and agent-based models. Ecol. Modell. 198, 115–126. doi:10.1016/j.ecolmodel.2006.04.023

Grimm, V., Berger, U., DeAngelis, D.L., Polhill, J.G., Giske, J., Railsback, S.F., 2010. The ODD protocol: A review and first update. Ecol. Modell. 221, 2760–2768. doi:10.1016/j.ecolmodel.2010.08.019

Hartig, F., Drechsler, M., 2010. Stay by thy neighbor? Social organization determines the efficiency of biodiversity markets with spatial incentives. Ecol. Complex. 7, 91–99. doi:10.1016/j.ecocom.2009.07.001

Jack, B.K., Kousky, C., Sims, K.R.E., 2008. Designing payments for ecosystem services:

Lessons from previous experience with incentive-based mechanisms. Proc. Natl. Acad. Sci. U. S. A. 105, 9465–70. doi:10.1073/pnas.0705503104

Kassam, A., Friedrich, T., 2010. Conservation Agriculture : Concepts , worldwide experience , and lessons for success of CA-based systems in the semi-arid Mediterranean environments. Options Méditerranéennes 11–51.

Kniveton, D., Smith, C., Wood, S., 2011. Agent-based model simulations of future changes in migration flows for Burkina Faso. Glob. Environ. Chang. 21, S34–S40. doi:10.1016/j.gloenvcha.2011.09.006

Kroeger, T., Casey, F., 2007. An assessment of market-based approaches to providing ecosystem services on agricultural lands. Ecol. Econ. 64, 321–332. doi:10.1016/j.ecolecon.2007.07.021

Langyintuo, A.S., Mekuria, M., 2008. Assessing the influence of neighborhood effects on the adoption of improved agricultural technologies in developing agriculture. African J. Agric. Resour. Econ. 2, 151–169.

Marenya, P., Smith, V.H., Nkonya, E., 2014. Relative Preferences for Soil Conservation Incentives among Smallholder Farmers: Evidence from Malawi. Am. J. Agric. Econ. 1–21. doi:10.1093/ajae/aat117

Marra, M., Pannell, D.J., Abadi Ghadim, A., 2003. The economics of risk, uncertainty and learning in the adoption of new agricultural technologies: where are we on the learning curve? Agric. Syst. 75, 215–234. doi:10.1016/S0308-521X(02)00066-5

Mondal, P., Basu, M., Bhadoria, P.B.S., 2011. Critical Review of Precision Agriculture Technologies and Its Scope of Adoption in India. Am. J. Exp. Agric. 1, 49–68.

Ndah, H.T., Schuler, J., Uthes, S., Zander, P., Traore, K., Gama, M.-S., Nyagumbo, I., Triomphe, B., Sieber, S., Corbeels, M., 2014. Adoption potential of conservation agriculture practices in sub-Saharan Africa: results from five case studies. Environ. Manage. 53, 620–35. doi:10.1007/s00267-013-0215-5

Palm, C., Blanco-Canqui, H., DeClerck, F., Gatere, L., Grace, P., 2014. Conservation agriculture and ecosystem services: An overview. Agric. Ecosyst. Environ. 187, 87–105. doi:10.1016/j.agee.2013.10.010

Pannell, D.J., Llewellyn, R.S., Corbeels, M., 2014. The farm-level economics of conservation agriculture for resource-poor farmers. Agric. Ecosyst. Environ. 187, 52–64. doi:10.1016/j.agee.2013.10.014

Pannell, D.J., Marshall, G.R., Barr, N., Curtis, a., Vanclay, F., Wilkinson, R., 2006.

Understanding and promoting adoption of conservation practices by rural landholders. Aust. J. Exp. Agric. 46, 1407. doi:10.1071/EA05037

Parkhurst, G.M., Shogren, J.F., 2007. Spatial incentives to coordinate contiguous habitat. Ecol. Econ. 64, 344–355. doi:10.1016/j.ecolecon.2007.07.009

Parkhurst, G.M., Shogren, J.F., 2008. Smart Subsidies for Conservation. Am. J. Agric. Econ. 90, 1192–1200. doi:10.1111/j.1467-8276.2008.01203.x

Parkhurst, G.M., Shogren, J.F., Bastian, C., Kivi, P., Donner, J., Smith, R.B.W., 2002. Agglomeration bonus: an incentive mechanism to reunite fragmented habitat for biodiversity conservation. Ecol. Econ. 41, 305–328. doi:10.1016/S0921-8009(02)00036-8

Pretty, J.N., Noble, A.D., Bossio, D., Dixon, J., Hine, R.E., Penning De Vries, F.W.T., Morison, J.I.L., 2006. Resource-conserving agriculture increases yields in developing countries. Environ. Sci. Technol. 40, 1114–1119.

Sci-kit Learn, 2015. 3.2.4.3.1. sklearn.ensemble.RandomForestClassifier [WWW Document].

Sutherland, L.-A., Gabriel, D., Hathaway-Jenkins, L., Pascual, U., Schmutz, U., Rigby, D., Godwin, R., Sait, S.M., Sakrabani, R., Kunin, W.E., Benton, T.G., Stagl, S., 2012. The "Neighbourhood Effect": A multidisciplinary assessment of the case for farmer co-ordination in agri-environmental programmes. Land use policy 29, 502–512. doi:10.1016/j.landusepol.2011.09.003

Tilman, D., Fargione, J., Wolff, B., Antonio, C.D., Dobson, A., Howarth, R., Schindler, D., Schlesinger, W.H., Simberloff, D., Swackhamer, D., 2001. Forecasting Agriculturally Driven Global Environmental Change 292, 281–285.

Vanclay, F., Lawrence, G., 1994. Farmer rationality and the adoption of environmentally sound practices; A critique of the assumptions of traditional agricultural extension. Eur. J. Agric. Educ. Ext. 1, 59–90. doi:10.1080/13892249485300061

Ward, P.S., Bell, A.R., Parkhurst, G.M., 2015. Heterogeneous Preferences and the Effects of Incentives in Promoting Conservation Agriculture in Malawi (No. 01440), IFPRI Discussion Paper Series. Washington, DC.

Wu, X., Murray, a. T., 2008. A new approach to quantifying spatial contiguity using graph theory and spatial interaction. Int. J. Geogr. Inf. Sci. 22, 387–407. doi:10.1080/13658810701405615

FIGURE CAPTIONS

Figure 1: Sample time path for yields from conventional and conservation land use, in the case of a farmer with technical efficiency of 0, soil quality of 0, with a mean rainfall of 250 units and standard deviation 90. Conservation farmers pay a big cost initially, but after 8 years on average both slightly outperform conventional farmers and have more resilient yields.

Figure 2 – Interaction Loop Expanded

Figure 3: Typical landscape (shown here a single instance with a base payment of 40, agglomeration payment of 60, penalty for cheating of 1000, and chance of being caught of 60%) during simulation run, at timesteps A) 0, B) 5, C) 20, and D) 30. Yellow plots are under conventional practice, green plots are under conservation practice, while red plots are cheating – maintaining conventional practice but claiming program benefits.

Figure 4: Proportion of total plot area adopting conservation practice, as function of total program spending on payments (X axis, along left axis within each panel) and the percentage of payments coming as agglomeration payments (Y axis, along right axis within each panel), for three different conditions of monitoring effectiveness (Low – per-area penalty of 1000 and chance of being caught 0.4; Medium – per-area penalty of 1500 and chance of being caught 0.6; High – per-area penalty of 2000 and chance of being caught 0.8). Surface is irregularly shaped as both x- and y-axis variables are modeled outcomes, not input variables. Spending on payments (X axis) is re-scaled to range from 0 to 1. The contours are equal area under conservation practice showing (especially evident in panels B and C) that the same area can occur with lower costs if a greater proportion of the payments are agglomeration ones.

Figure 5: Proportion of total plot area under conservation practice, as function of spending on payments and the likelihood of being caught, for several subsets of data: A) all model cases where agglomeration payments are 0 (with base payments spanning 0 to 100), B) all model cases where base payments are 0 (with agglomeration payments spanning 0 to 100), C) all model cases where base payments are 40, and D) all model cases where base payments are 80. Color (from blue through red) represents proportion of total area under conservation practice; contours of same color have same level of area under conservation practice. Black dashed line represents the maximum program spending in model cases with only base payments. The "L-shaped" contours in B-D (absent in A) indicate that different combinations of monitoring effort (proxied by the likelihood of being caught) and agglomeration payments may lead to the same adoption area. For example, the maximum area using base payments only (at a cost of 0.375 and monitoring effort of 0.4) in A, can be achieved with base and agglomeration payments at half the cost (~0.18) at the same monitoring effort in B, or at a higher cost with lower monitoring,

or *vice versa*. Surface is irregularly shaped as y-axis variable is a modeled outcome, not an input variable.

Figure 6 – Relative importances of Monte Carlo variables in predicting modeled A) area under adoption and B) program spending on payments, estimated using 50 different samples of 25,000 data points from our full dataset with a random forest classifier. Standard deviations shown in horizontal blue bars. Shortnames shown here are defined in Table 1.

Table 1 – Model Parameter Settings for Experiment

Modeling Parameters - Varied in Monte Carlo Analysis	Shortname (in Figure 6)	Minimum	Maximum	Units
Number of interaction steps per decision time step	*interactionsPerTimeStep*	5	9	/year
Rainfall Mean	*climateMean*	250	500	mm/year*
Rainfall SD	*climateSD*	50	100	mm/year*
Number of time steps / years considered in estimating expected utility	*numYearsForward*	15	25	year
Transaction cost, switch to conservation practice	*switchEncouragedCost*	50	300	dollar*
Transaction cost, switch to conventional practice	*switchStatusQuoCost*	50	300	dollar*
Average link strength for network	*connectionScalingFactor*	0.3	0.8	-
Plot Radius Mean	*propertySizeMean*	50	80	m*
Plot Radius SD	*propertySizeSD*	20	80	m*
Network link Kernel Radius	*networkCorrelationDistance*	500	2000	m*
Plot Soil Quality SD	*agentSoilSD*	0.05	0.1	-
Farm Discount Rate Mean	*agentDiscountMean*	0.02	0.1	-
Farm Discount Rate SD	*agentDiscountSD*	0.01	0.05	-
Farm Risk Preference Mean	*agentRValueMean*	0.5	1.25	-
Farm Risk Preference SD	*agentRVAlueSD*	0.1	0.3	-
Farm Participation in Interaction Step Mean	*agentInteractChanceMean*		0.9	-
Farm Participation in Interaction Step SD	*agentInteractChanceSD*	0.1	0.2	-
Number of random draws to calculate expected utility	*numDraws*	10	20	-
Perceived likelihood of a given neighbor adopting, for use when estimating expected utility	*chanceAdopt*	0.3	0.7	-
Number of plots per farm Mean	*meanPlotsFarm*	2	4	-
Number of plots per farm SD	*sdPlotsFarm*	0	3	-
Number of options considered Mean	*meanNumCombos*	6	14	-
Number of options considered SD	*sdNumCombos*	1	3	-
Probability of leaving landscape	*probLeave*	0.01	0.04	-
Farm technical efficiency SD	*sdEff*	0.05	0.1	-
Probability that a farm honors an offer to a neighbor that cheats	*probGiveCheat*	0	1	-

Modeling Parameters - Fixed in all simulations	Value	Units
Fraction of landscape filled	0.6	-
Landscape X	3000	m*
Landscape Y	3000	m*
Years in simulation	30	year
Number of sponsored pilot study members	10	farm
Start year for early adoption pilot	5	-
End year for early adoption pilot	8	-
Length of subsidy	6	years

Modeling Parameters - Varied in Policy Sweep	Minimum	Increment	Maximum	Units
Base Payment	0	20	100	dollar/year*
Agglomeration Payment	0	20	100	dollar/year*
Penalty for Cheating	1000	500	2000	dollar/year*
Random monitoring Probability	0.2	0.2	0.8	-

Units are provided where appropriate, but we emphasize that in an abstract model such as this, they are useful only as a guide

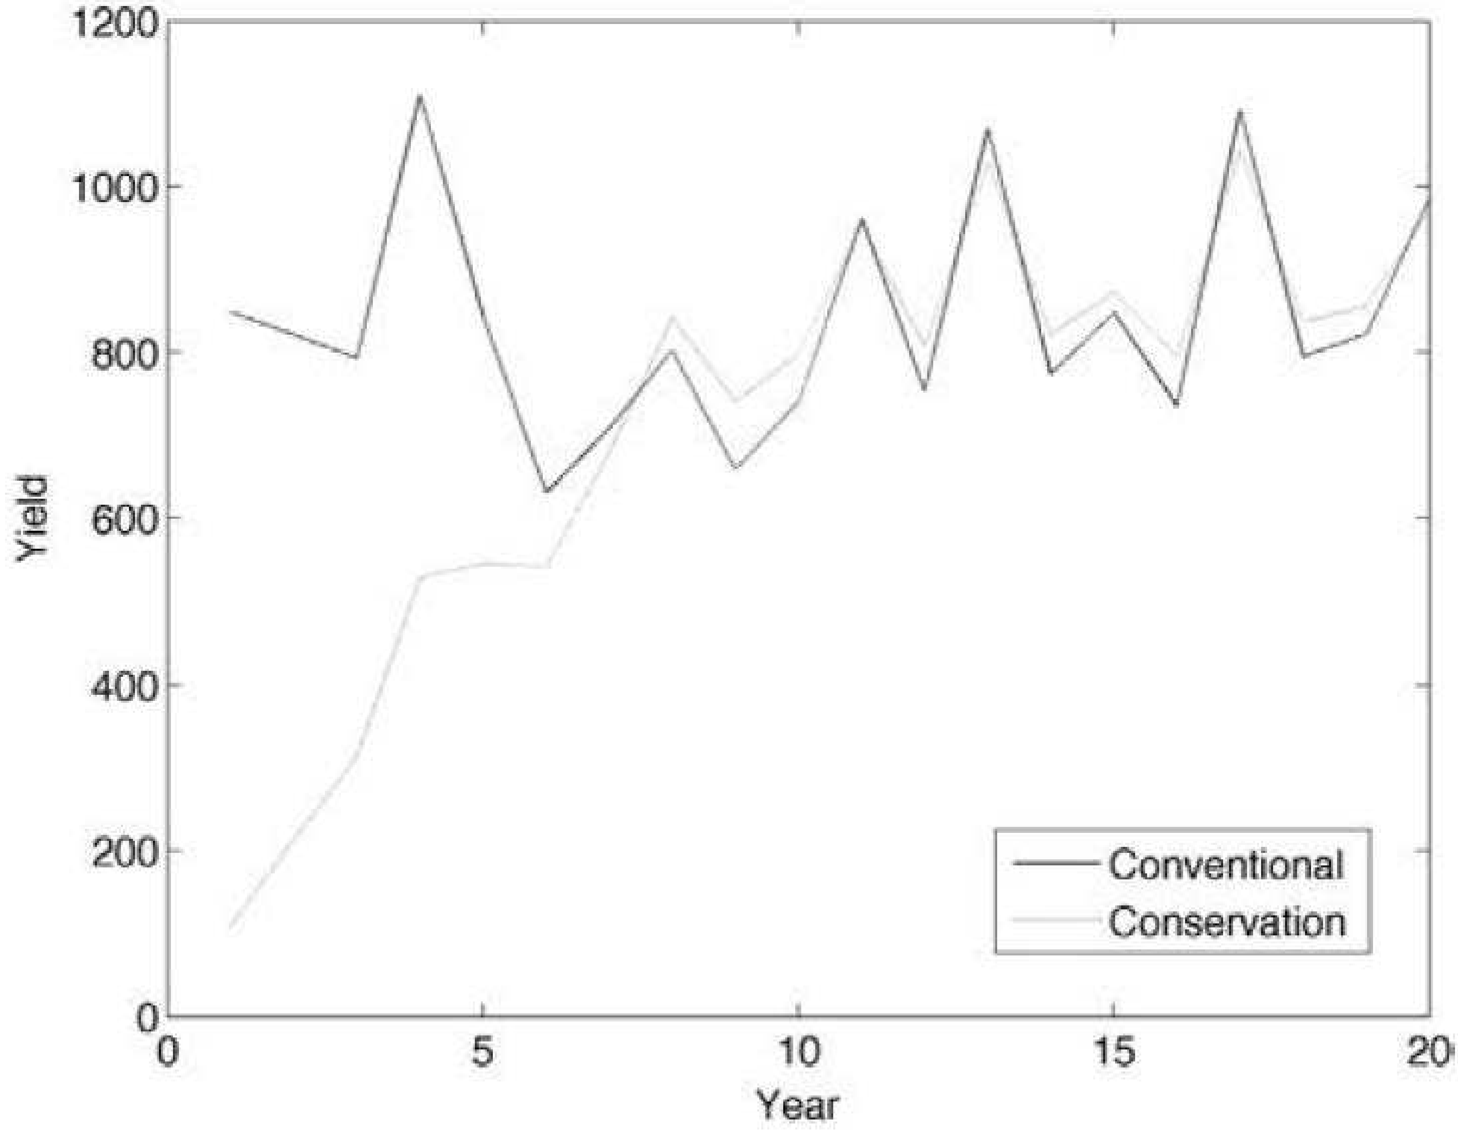

1200
1000
800
600
400
200
0
Yield
0
5
10
15
20
Year
Conventional
Conservation

Interaction Loop (Expanded)

Nested loop of all farms j within interaction step k

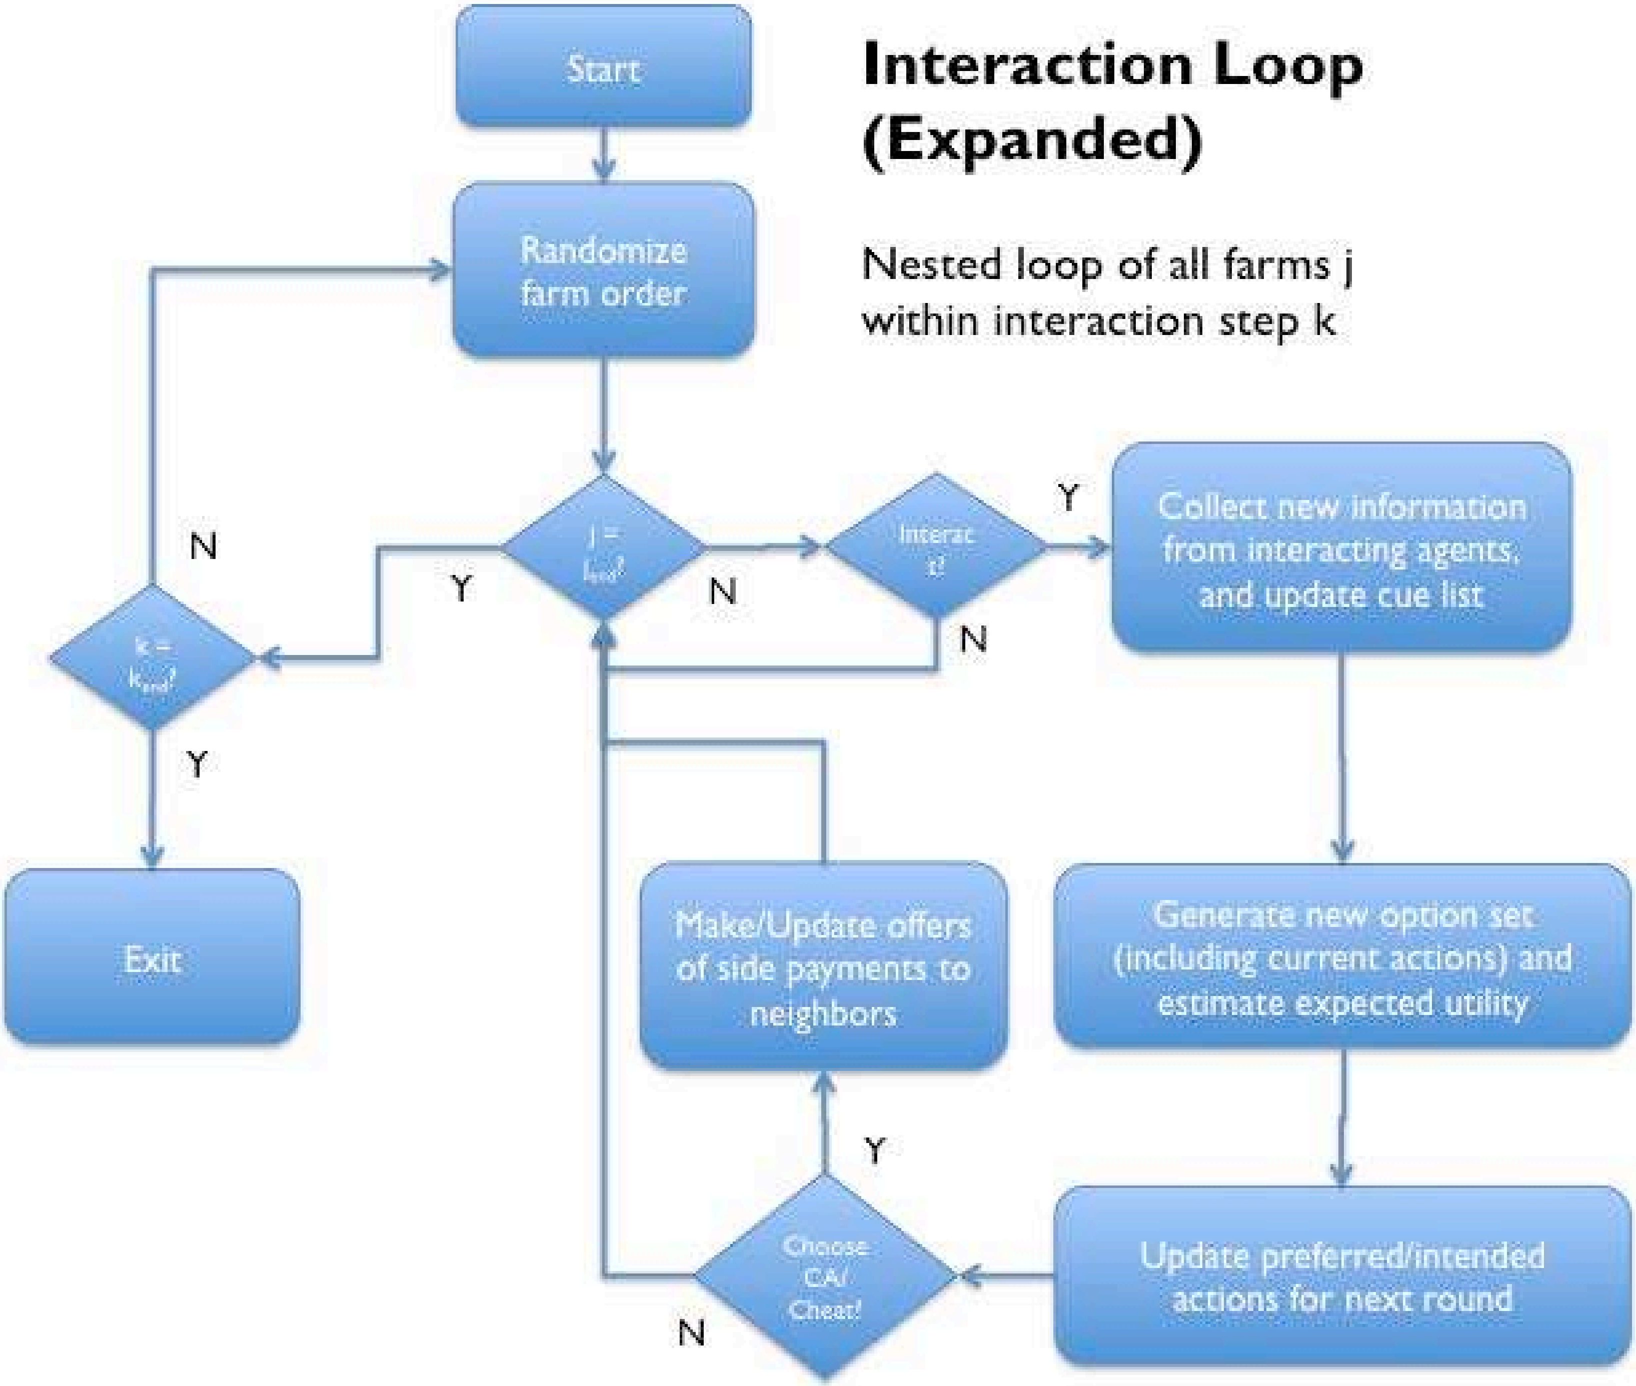

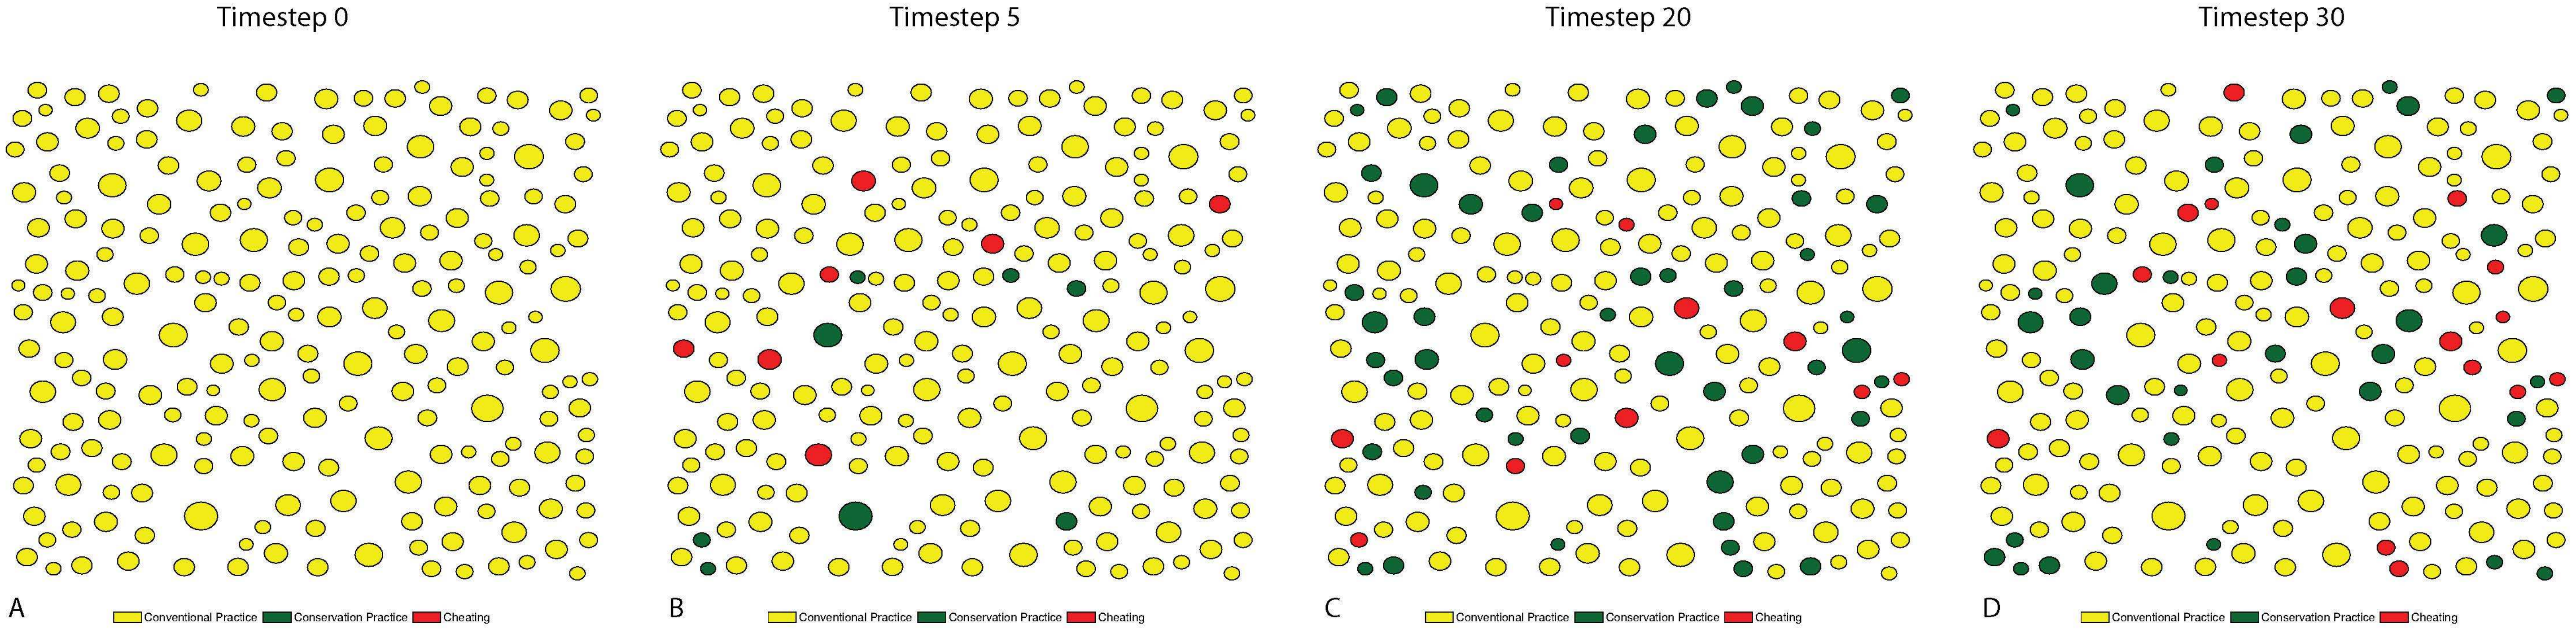

Timestep 0
Timestep 5
Timestep 20
Timestep 30
A
Conventional Practice
Conservation Practice
Cheating
B
Conventional Practice
Conservation Practice
Cheating
C
Conventional Practice
Conservation Practice
Cheating
D
Conventional Practice
Conservation Practice
Cheating

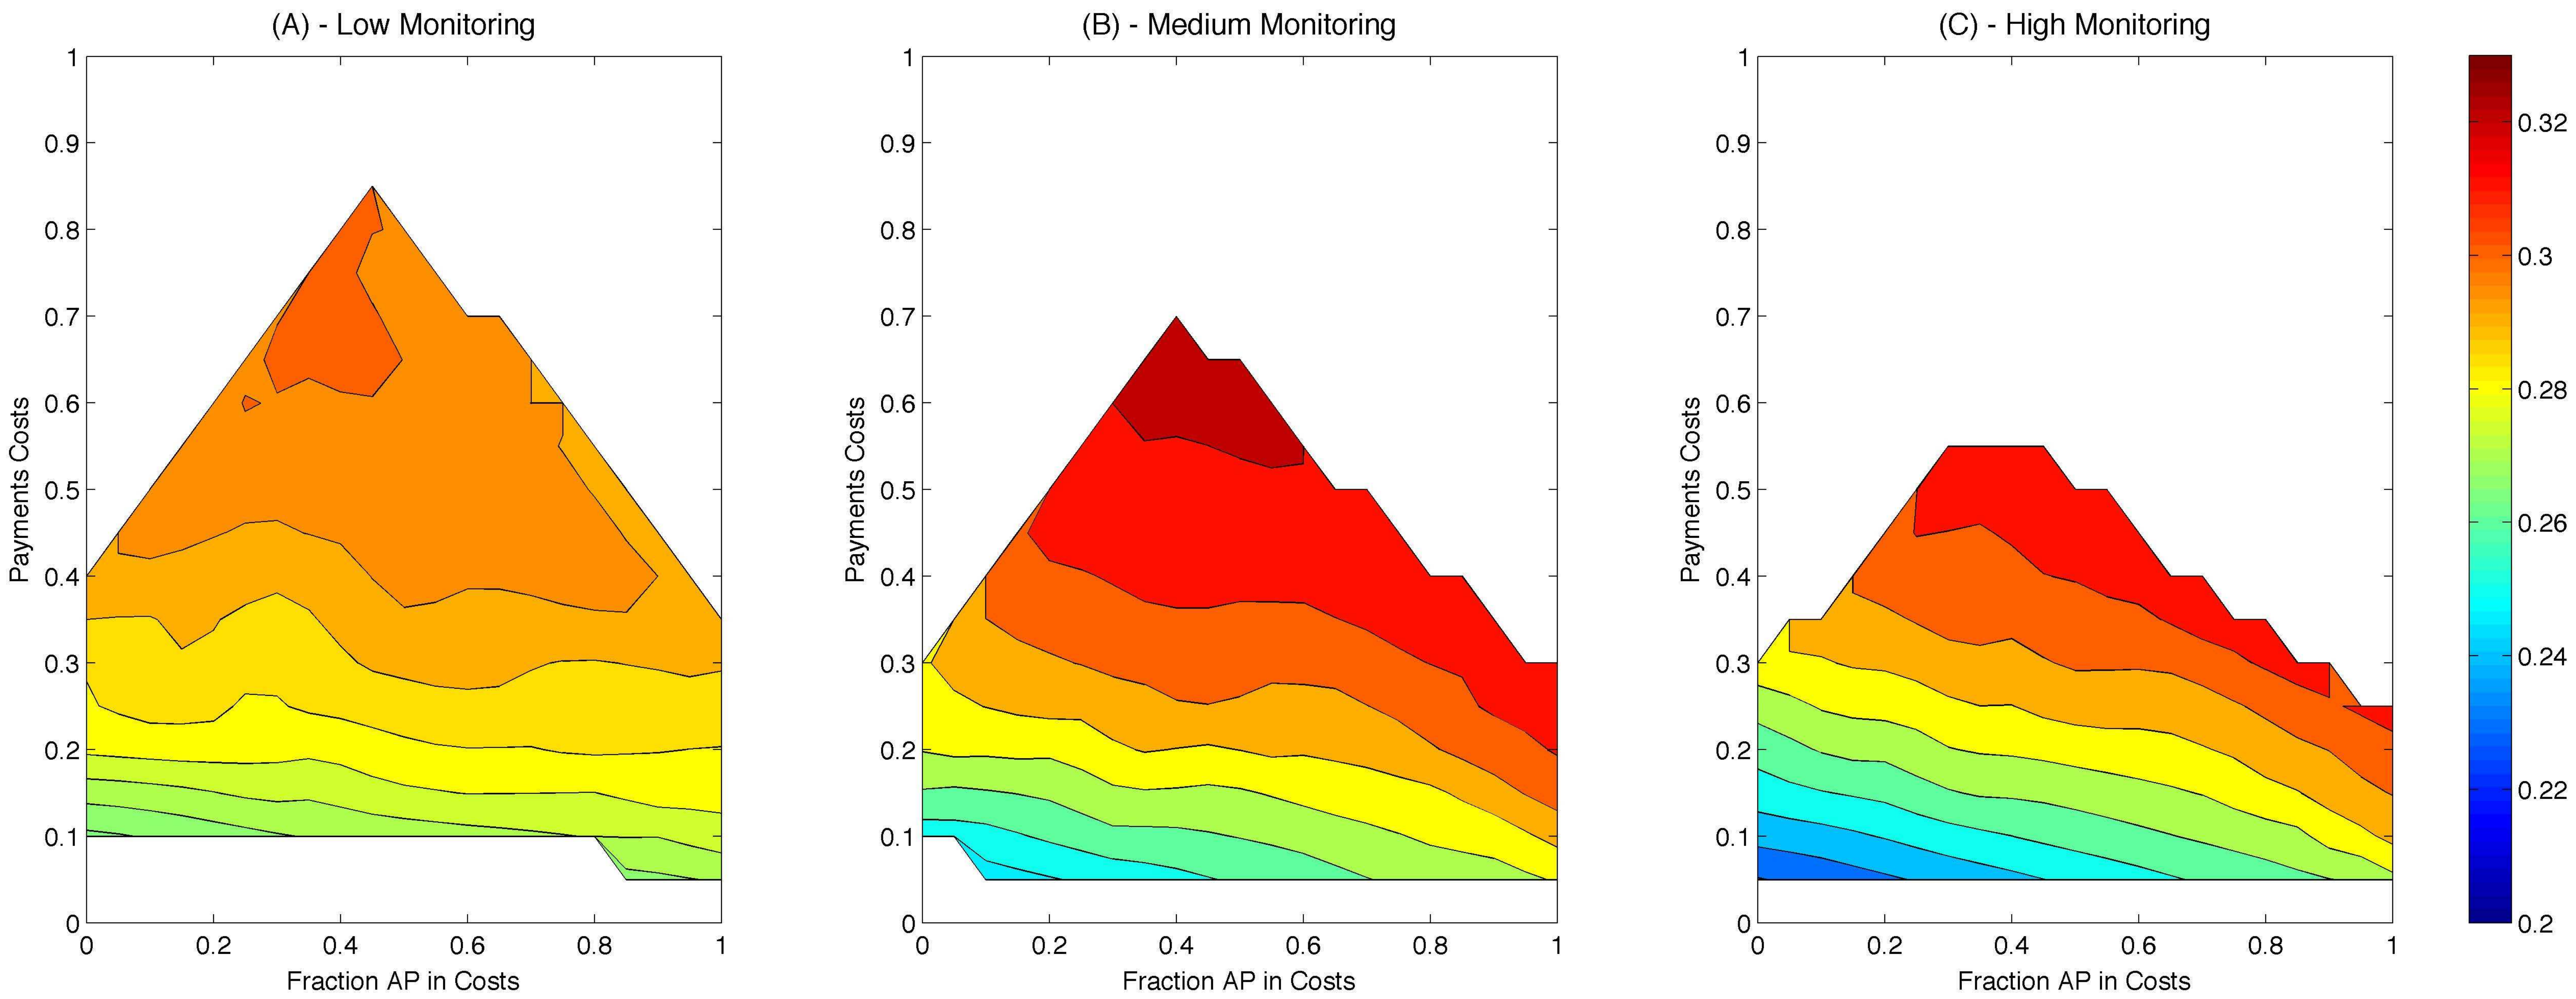

(A) - Low Monitoring
(B) - Medium Monitoring
(C) - High Monitoring
Payments Costs
Fraction AP in Costs
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0.2
0.22
0.24
0.26
0.28
0.3
0.32

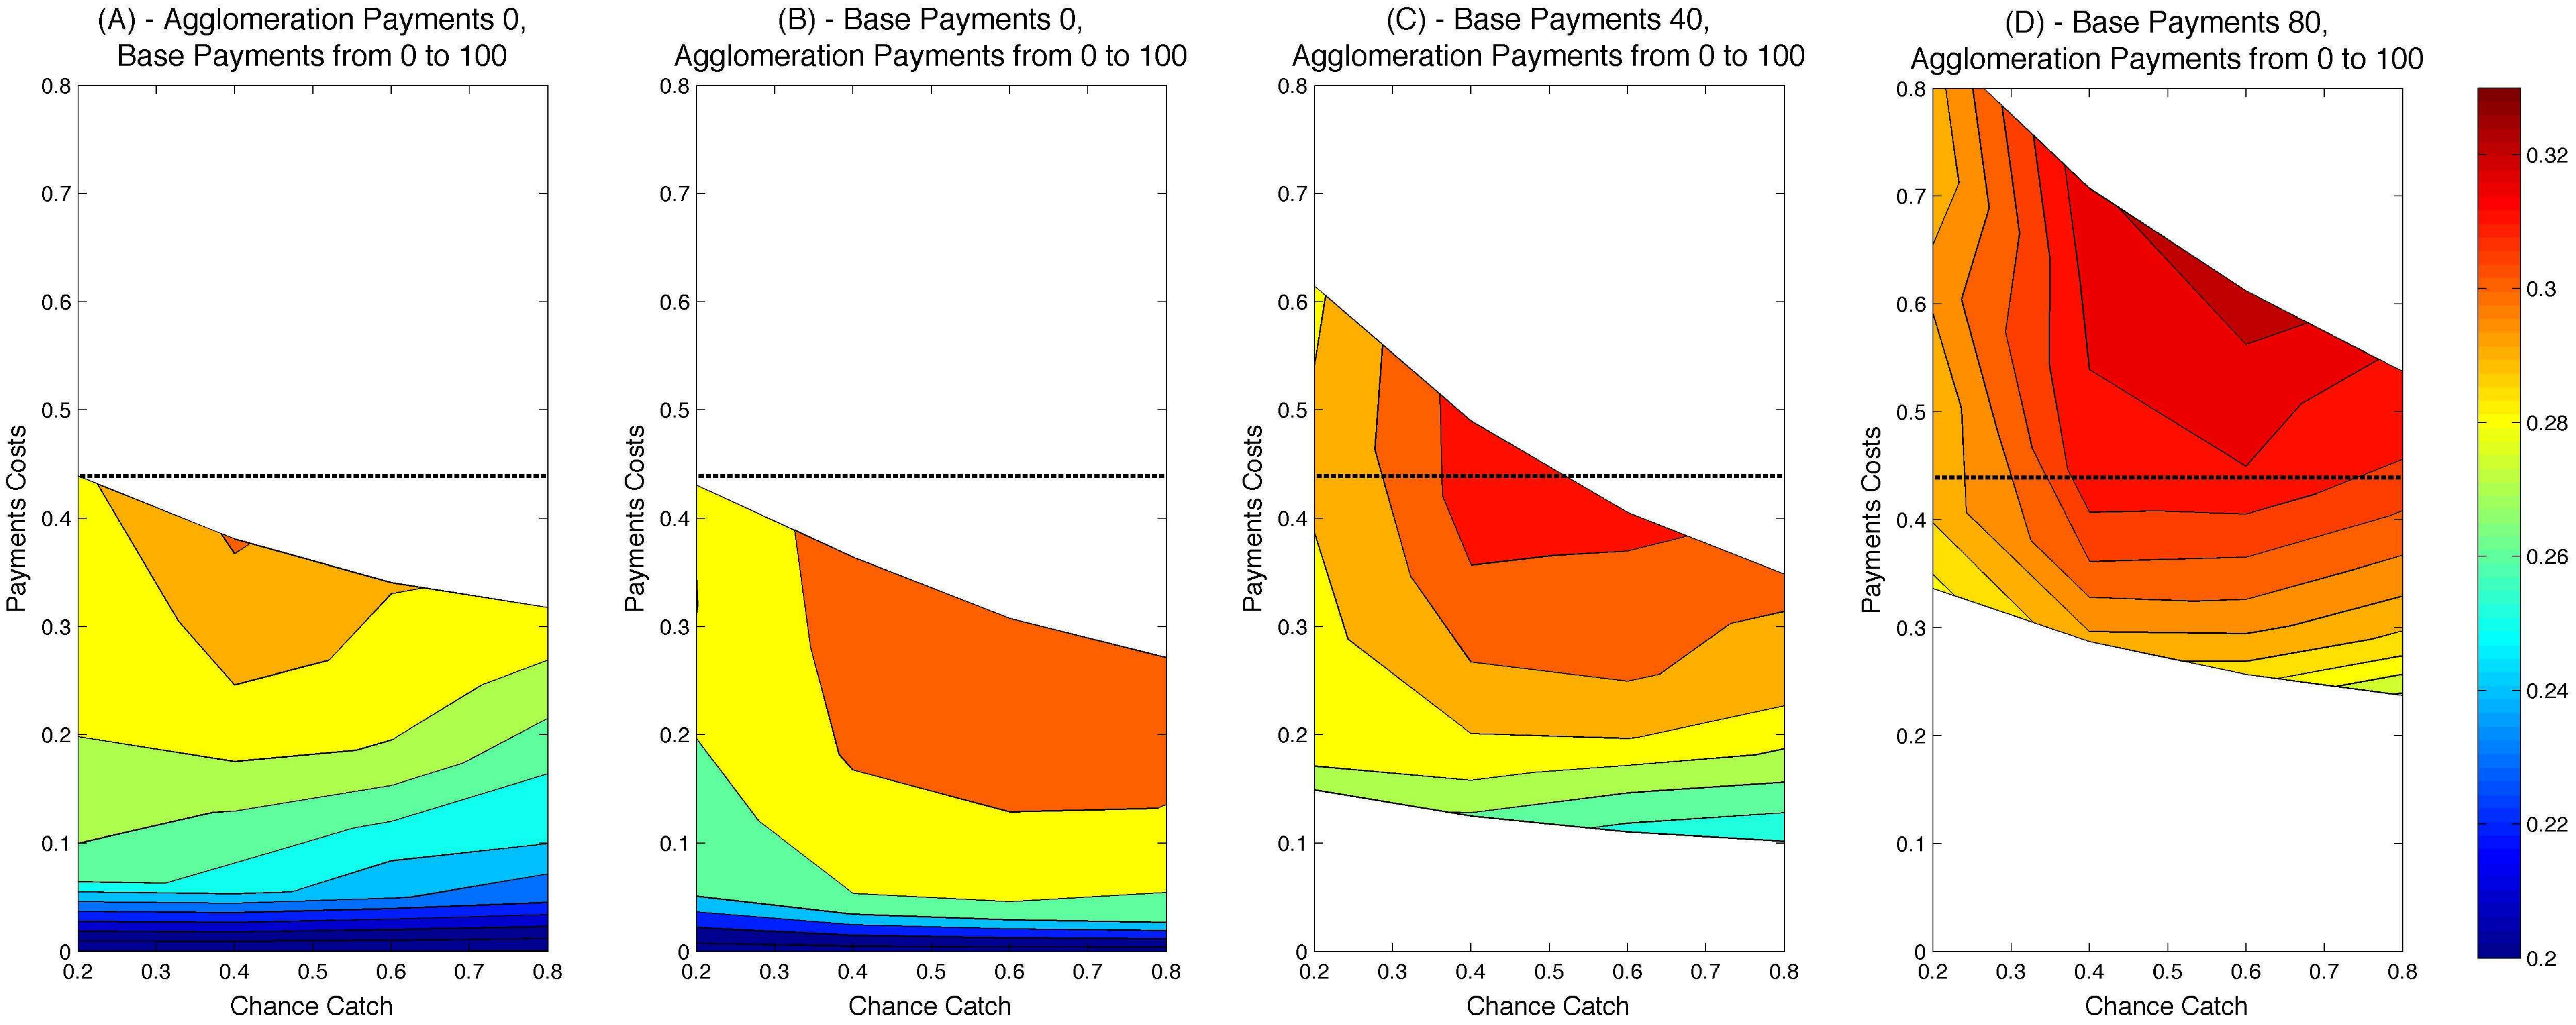
(A) - Agglomeration Payments 0,
Base Payments from 0 to 100
(B) - Base Payments 0,
Agglomeration Payments from 0 to 100
(C) - Base Payments 40,
Agglomeration Payments from 0 to 100
(D) - Base Payments 80,
Agglomeration Payments from 0 to 100
Payments Costs
Chance Catch
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.2
0.22
0.24
0.26
0.28
0.3
0.32

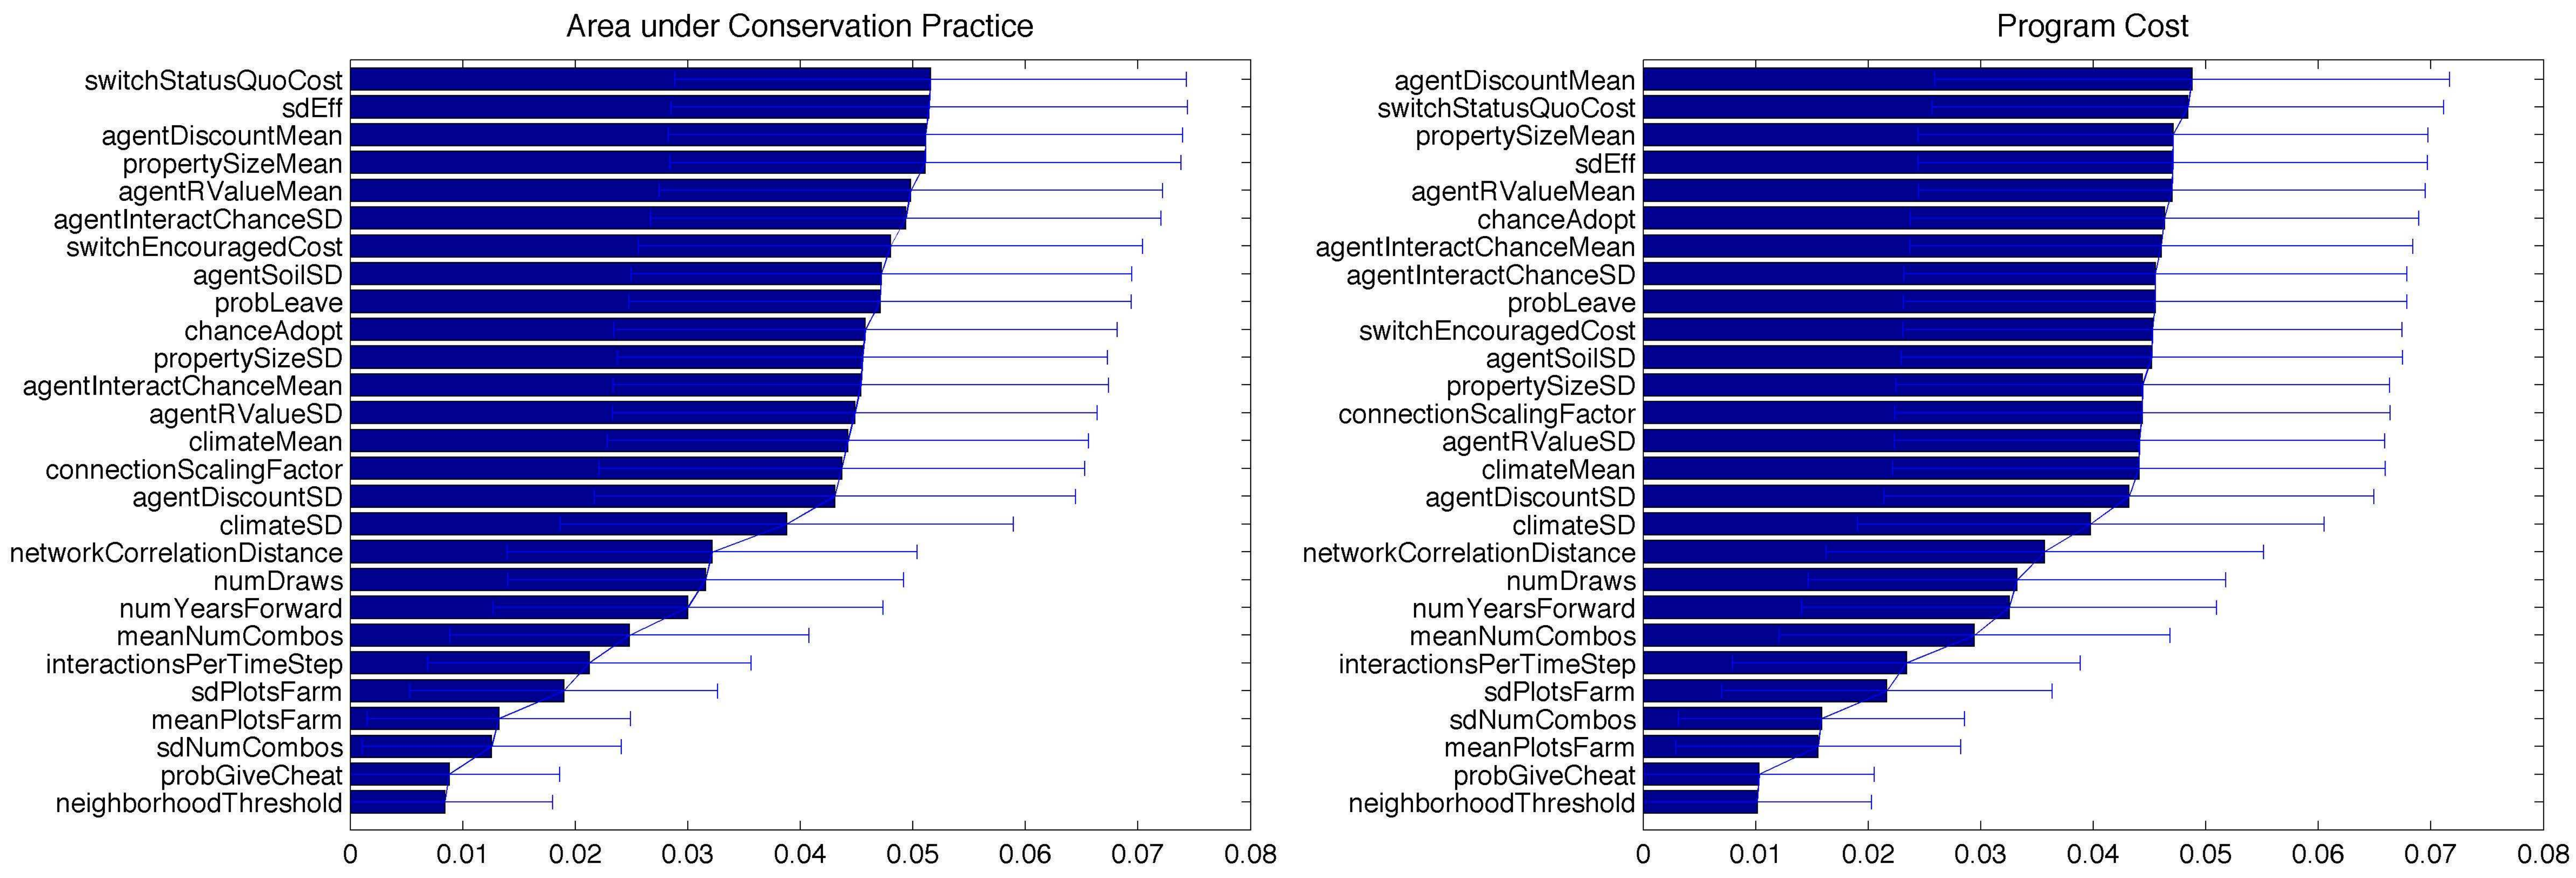

Area under Conservation Practice
switchStatusQuoCost
sdEff
agentDiscountMean
propertySizeMean
agentRValueMean
agentInteractChanceSD
switchEncouragedCost
agentSoilSD
probLeave
chanceAdopt
propertySizeSD
agentInteractChanceMean
agentRValueSD
climateMean
connectionScalingFactor
agentDiscountSD
climateSD
networkCorrelationDistance
numDraws
numYearsForward
meanNumCombos
interactionsPerTimeStep
sdPlotsFarm
meanPlotsFarm
sdNumCombos
probGiveCheat
neighborhoodThreshold
0
0.01
0.02
0.03
0.04
0.05
0.06
0.07
0.08
Program Cost
agentDiscountMean
switchStatusQuoCost
propertySizeMean
sdEff
agentRValueMean
chanceAdopt
agentInteractChanceMean
agentInteractChanceSD
probLeave
switchEncouragedCost
agentSoilSD
propertySizeSD
connectionScalingFactor
agentRValueSD
climateMean
agentDiscountSD
climateSD
networkCorrelationDistance
numDraws
numYearsForward
meanNumCombos
interactionsPerTimeStep
sdPlotsFarm
sdNumCombos
meanPlotsFarm
probGiveCheat
neighborhoodThreshold
0
0.01
0.02
0.03
0.04
0.05
0.06
0.07
0.08

Model Description

We describe our agent-based model using the ODD (Overview, Design concepts, Details) protocol for describing individual- and agent-based models (Grimm et al., 2006, 2010).

1. Purpose

To assess the performance of agglomeration payments alongside conventional subsidies in encouraging conservation practice in developing country agricultural contexts, across a broad range of farmer and environmental characteristics.

2. Entities, state variables, and scales

There are four entities used in this model – plots, farms, environment, and networks. Additionally, there are two time scales – the interaction time scale (governing interactions among farmers) and the decision time scale (governing decisions made and actions taken on plots) and one spatial scale (the farm landscape).

Plots

Plots represent the spaces upon which farms cultivate crops. They are described by: a unique ID, a location in two-dimensional space and a size, an owner farm, the current state (one of three management choices), the number of periods since the state was changed, and a single land quality parameter. The state of each plot is either i) managed using conventional practice, ii) registered for the payment program and managed using conservation practice, or iii) 'cheating' – i.e., registered for the payment program but still managed using conventional practice.

Farms

Farms represent farming households as a single decision-making body. They are described by: a unique ID, a list of plots owned, a single variable for accumulated wealth, a list of network links to other farms (described in detail below in *Collectives*), a single variable for technical efficiency, a risk aversion coefficient, a discount rate, and a memory of their own past experiences on their plots as well as those experiences shared with them by other farmers with whom they have interacted. Additionally, they are described by state variables specific to model processes, including a list of offers outstanding to neighboring farmers to encourage adoption of the conservation practice (see *Process Scheduling*), the maximum offer they would make in such a situation, the fraction of experiences in their memory they are able to recall when deciding among options (see *Process Scheduling*), the number of options they consider at a time, the probability of participating in a round of interactions with other farmers (see *Time Scales*), a flag for whether they were selected as part of the initial pilot program in the model (see *Initialization*), and a flag for whether offers they've made to neighbors would be honored if the neighbor cheated (i.e., registered in the conservation practices program but did not follow through; see *Process Scheduling*).

Environment

The environment in which farms (and plots) are embedded is described by a small set of physical and social parameters: the size of the two-dimensional landscape, the fraction of the landscape to be filled with plots, a threshold radius around a plot for what is considered a neighbor, a normal distribution for rainfall from which actual rainfall is drawn for each plot, a transaction cost for switching between the conventional practice and the conservation practice, a transaction cost for switching from the conservation practice back to the conventional practice, a time path for conventional subsidies being offered for registration in the conservation practice, a time path for agglomeration payments being offered for registration in the conservation practice, a penalty for being caught cheating and a likelihood of being caught, the number of farms initially selected to participate in the pilot program (see *Initialization*), and the start and end of this initial pilot.

Networks

Farms are connected to other farms by network links; the link strength (from 0 to 1) indicates the probability that a farm will share information about past experience with the linked farm in a given interaction, and is directional – the link strength from farm *i* to farm *j* is independent of the link strength from farm *j* to farm *i*. Link strengths are assigned randomly between plots, and then scaled down using a linear kernel outside the neighborhood radius (i.e., the average strength of links decays outside of local neighborhoods); when plots are assigned to farms, the strongest link between any plots owned by farm *i* and any plots owned by farm *j* is retained as the link strength between those two farms. This allows link strength to depend on physical neighborhoods even when farms are composed of plots spaced far apart, common in smallholder environments.

Scales

Simulations in this study take place on a 3km x 3km square space, with plots located continuously across the space (not gridded). The decision time step represents one year, as though in an environment with a single agricultural season per year (common to Sub-Saharan Africa), and simulations are carried through for 30 years. The interaction time step does not map to an invariantly regular pace of interaction among farms, but captures the reality that farmers may have a range of interactions with other farmers in between making decisions on their farms. The number of interaction time steps between each decision time step varies across our experiment (see *Experiment Setup*) from 5 to 9; each farm selects probabilistically to engage other farms at all in each interaction step, and in each interaction information is shared probabilistically based on the strength of the network link. This allows a distribution in the number of interactions shared among farms between decisions, rather than a spuriously invariant number of interactions.

3. Process overview and scheduling

The main loop of the model consists of three inner loops that update farm yields (Loop 1), information (Loop 2), and actions (Loop 3) once for each decision time step (Figure 1).

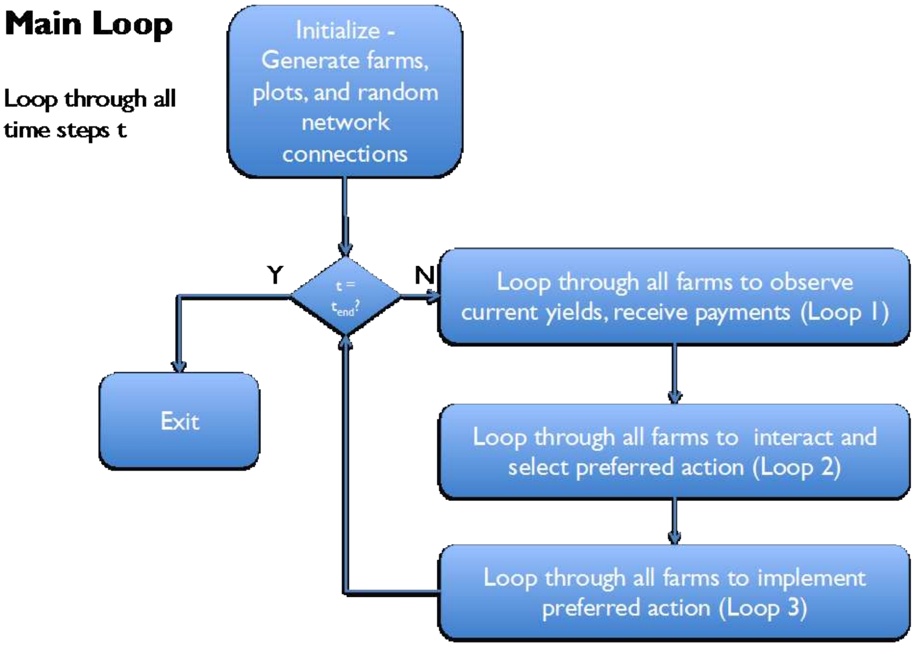

Figure 1 – Main Model Loop

Loop 1 generates rainfall for each plot in the landscape, then cycles through each farm in the landscape, evaluating the yields, program payments, and penalties accruing to each plot in the farm (Figure 2). Plots that are 'cheating' have a finite probability of being caught, in which case a penalty of fixed and known value accrues to the farm. Farms add current observations of net yields (in the case of the 'Cheat' choice, penalties assessed are subtracted from yields) to their list of experience. At the end of the main loop, farms leave the landscape with some probability, and their farms are taken over by new households (farm parameters are re-estimated but farm plots do not change). Updating of farms is asynchronous.

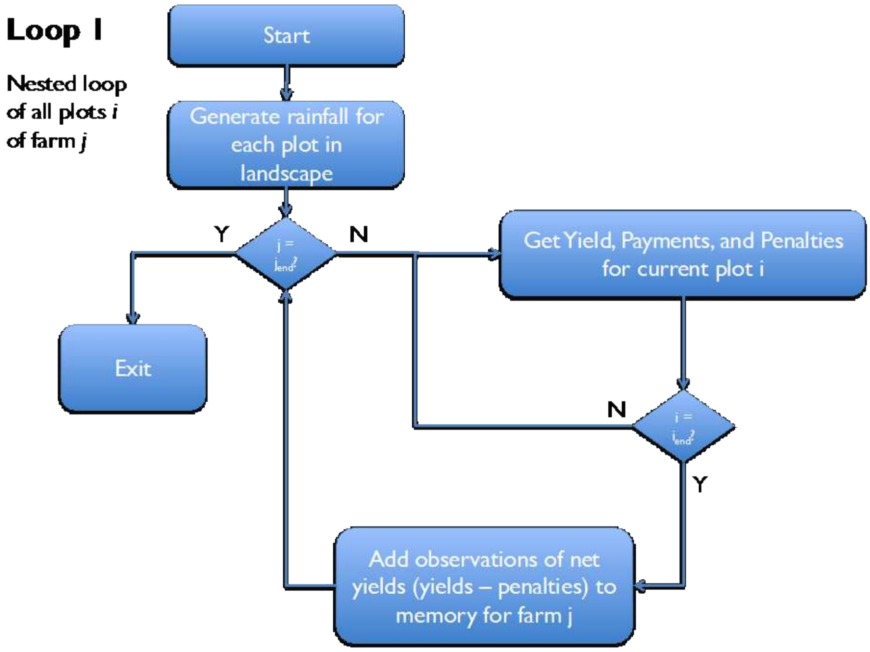

Figure 2 – Loop 1 (Yields)

Loop 2 captures the interaction time steps that occur between decision time steps (Figure 3). For each of k interaction timesteps, the ordering of farms is randomized, and the model loops through each farm. If the farm participates in this interaction time step (determined by random draw), it will then exchange information with other farms in its network, probabilistically depending on the strength of the network link with each farm. (Note, even if a farm does not 'speak' in a particular interaction timestep, it may still be 'spoken to' by other farms that do participate). Once information is collected, farms update a 'cue list', which stores a similarity measure (scaled from 0 to 1) of the other farms relative to themselves. This similarity is calculated using the mean deviation in per-area yield when farms were making the same choices (conventional, conservation, or cheat) in the same year; cues are scaled so that the most similar farms have a cue of 1 and the least similar have a cue of 0. Again, updating of farms is asynchronous.

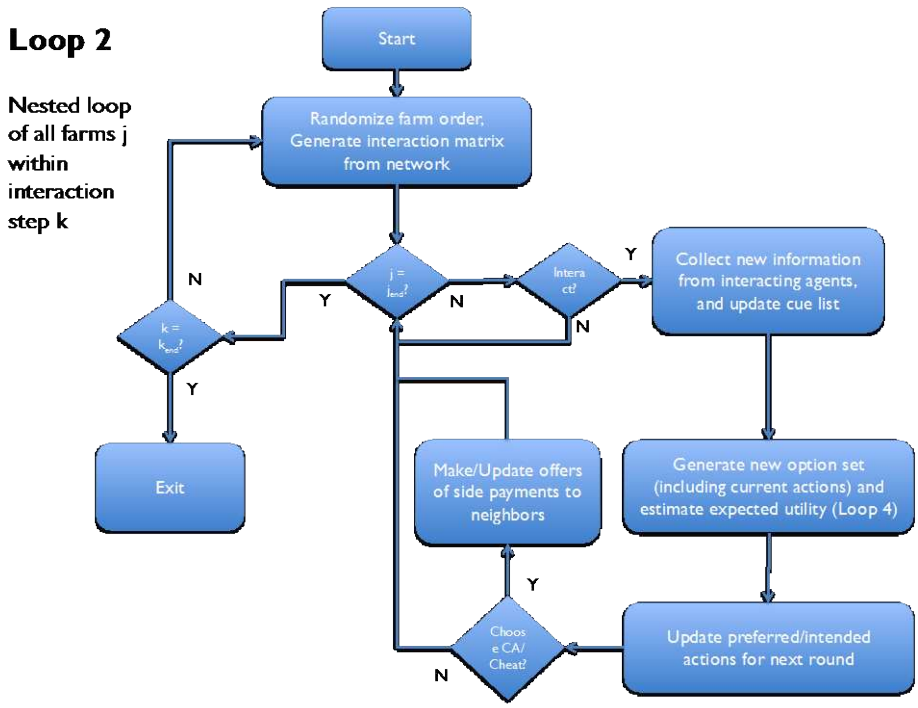

Figure 3 – Loop 2 (Information)

Once the farm has updated its information, it then reconsiders its planned actions for the upcoming decision time step. A set of n different options, including the current land use, is generated, and the farm generates an estimate of the expected utility of each option (see Loop 4 in *Submodels*). There are three different actions possible for a plot, so that a farm composed of *w* plots has 3^w different configurations possible, only a small subset of which are considered in a particular turn. Based on the outcome of this evaluation, farms update their planned actions for the upcoming decision time step. If the planned or current actions include doing conservation practice, or cheating, farms also provide an 'encouragement' or side payment to neighboring farms to encourage them to also adopt ''(or cheat) (see *Submodels*).

After all interactions have occurred, Loop 3 cycles through all farms to update their land use and, if applicable, pay transaction costs for changing land use and accept any side payments from farms (Figure 4). Updating of farms is asynchronous.

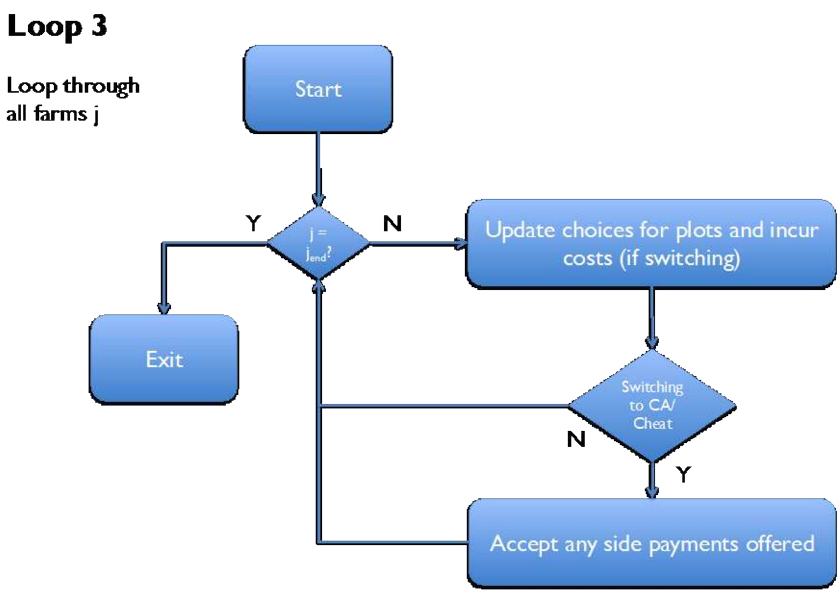

Figure 4 – Loop 3 (Actions)

4. Design concepts

Basic principles

This model brings together theories of bounded rationality (Kahneman, 2003; Rubenstein, 1998), risk aversion (G. Feder, 1980; Tanaka, Camerer, & Nguyen, 2010), and expected utility (Gershon Feder, Just, & Zilberman, 1985) to construct a decision model for rural smallholders in developing countries. In turn, this decision model allows farms to collect information, trial the conservation practices on their farms, and eventually they may adopt at greater degrees – key principles of adoption theory (Ghadim & Pannell, 1999; Pannell et al., 2006). Our sub-model for production under conventional and conservation practices embeds understanding that conservation practices can lead to improved yields, but only after some lag period (FAO, 2012; Pretty et al., 2006).

Emergence

Landscape-scale outcomes of adoption level, as well as contiguity of adoption, and program cost emerge out of the decisions of individual farms to adopt. Adoption can be understood in this model as including early, majority, and late adopters, with the model also simulating disadoption.

Adaptation

Farms change their land-use practices at the plot level, across three different options (conventional practice, conservation practice, and cheating – see *Sub-models*), using information from their own experience and the experience of others in their networks to identify options that maximize expected utility.

Objectives

Farms act to maximize expected utility from farm production and program payments.

Learning

Farms exchange observations on crop performance with other farmers.

Prediction

Using observations from other farms as well as their own experience, farms estimate the expected utility of different land-use options (see *Sub*-models).

Sensing
Farms are able to interact with other farms to exchange information about past crop experiences; additionally, farmers are aware of program conditions (value and length of conventional and agglomeration payments) as well as how their immediate neighbors are value the different options they have considered. Farms are also aware when neighbors have made offers to them that are conditional on adopting conservation practice (or cheating). These are the only environmental variables used in the estimation of expected utility.

Interaction
Farms interact with other farms probabilistically during each interaction time step; when an interaction occurs, farms share their list of past experiences in cropping with each other.

Stochasticity
Stochasticity is employed in the model in the following ways:

- Initial sizes and locations of plots are drawn randomly, and are assigned randomly to farms
- Farm characteristics, including discount rates and risk coefficients, are drawn randomly
- Rainfall for each plot is drawn randomly in each decision time step
- The set of possible options for consideration in each interaction time step is drawn randomly
- The observation data used to estimate expected utility in each draw, for each option is selected randomly
- The time path along which neighbors are imagined to adopt, for each draw of each option, is generated randomly
- The estimated cost of nudging neighbors to adopt is drawn randomly
- The initial offer made to neighbors to encourage them to adopt, as well as the amount to update offers, is drawn randomly
- The ordering with which farms go through the interaction time step is random
- Whether farms interact with each other in an interaction time step is drawn probabilistically
- Whether farms leave the landscape in a given turn is drawn randomly

Collectives
Farms are linked to other farms in the landscape via links whose strength scales from 0 through 1; a 1 indicates 100% probability that the two farms will interact to share information in a given interaction timestep.

Observation
All individual farm data are stored in model output; our analyses make use of aggregate landscape-scale outcomes of i) program costs and ii) area under conservation practice, over time.

5. Initialization

At initialization, a landscape of plots is randomly generated, with fraction f_{fill} of the landscape made up of plots with radii drawn from [$\mu_{r,plot}$, $\sigma_{r,plot}$]. Plots are randomly assigned into farms of [$\mu_{num\ plots}$, $\sigma_{num\ plots}$] each, and all other characteristics of the farm – risk behaviors, network link strengths, etc. (Table 1) are generated and assigned to the farms. The model proceeds up to timestep $t_{pilot\ start}$ with all plots using conventional practice. From timestep $t_{pilot\ start}$ through $t_{pilot\ end}$, a subset of n_{pilot} farms are selected to 'participate' in a pilot project of the conservation practice. Participation of these farms is achieved by manipulating the perceived utility derived from conservation practice, such that they choose an option for their farm that includes either i) conservation practice or ii) 'cheating' on at least one of their plots; the perceived utility is held high for this option for the duration of the pilot so that these farmers observe yields from conservation practice and cheating and store them in memory. Through this mechanism, knowledge about the conservation practice as well as cheating is inoculated into the landscape.

As outlined in Table 1, some model parameters are fixed across all simulations presented in the current study, while others vary. Specifically, parameters for the policy variables in the table are varied systematically over a fixed parameter range, with environmental and farm parameters drawn randomly from a range and then held fixed over one complete sweep of simulation runs through the policy parameter values. In this abstract model, parameter values are not meant to be pegged to particular literature values, and are largely arbitrary.

6. Input data

The model does not use input data to represent time-varying processes.

7. Submodels

Land use choices

Plots may take one of three states – i) conventional practice, ii) conservation practice, or iii) 'cheating' – claiming to do conservation practice in order to receive program payments but actually doing conventional practice. Farms do not choose the state of each plot directly; rather, they choose among different portfolios in a randomly generated set of options, where each portfolio represents a random configuration of their plots with these three states. The actual production functions for these three states are as follows:

- $Y_{conventional} = (600 + 1 \cdot Rainfall)(1 + 1.2 \cdot Soil)(1 + 1 \cdot Efficiency)$
- $Y_{conservation} = \frac{(\min(8, Seasons))}{8}(700 + 0.7 \cdot Rainfall)(1 + 1 \cdot Soil)(1 + 1.2 \cdot Efficiency)$
- $Y_{cheating} = (600 + 1 \cdot Rainfall)(1 + 1.2 \cdot Soil)(1 + 1 \cdot Efficiency) - (rand() < p_{catch}) \cdot Penalty$

These production functions are not tightly calibrated to actual production functions, but have the following structural characteristics:

- Yields for conservation practice are similar to conventional practice, but are i) slightly higher, ii) slightly less sensitive to the soil quality (plot specific) and iii) slightly more sensitive to the technical efficiency (farm specific)

- Yields for conservation practice are only fully realized after continuous adoption for many seasons

- Yields for cheating are the same as those for conventional practice, though there is a penalty that may be assessed if the cheating is caught (determined probabilistically)

These characteristics embed the notions of conservation agricultural practice reducing sensitivity of the yield to the environment, but requiring more from the farmer.

Expected Utility

The calculated expected utility for a plot-use portfolio is the measure by which farms compare different portfolios and select the use of their plots for the upcoming turn. For a portfolio *m* evaluated by farm *j*, the expected utility is:

$$\overline{U_{J,m}} = \frac{\sum_{w=1}^{n_{draws}} c_w \sum_{i=1}^{n_{years}} \left(\frac{P_{i,w,m}{}^{r_j}}{r_j}\right)^{-(1+d_j)}}{\sum_{w=1}^{n_{draws}} c_w}$$

where n_{draws} is the number of independent estimates of the utility time path made by the farm, n_{years} is the length in time periods of each time path considered, $P_{i,w,m}$ is the net value in year *i* expected to be earned from portfolio *m* using estimates of past yields from the farm(s) consulted for time path *w*, and c_w is the weight applied to the time path *w* based on the similarity in past performance of the farm(s) consulted for time path *w* with farm *j*. Similarity is calculated as the mean-square deviation between farms in per-area yields for the same action (conventional, conservation, or cheat) undertaken in the same year; the weights c_w are scaled so that the most similar farms have a weight of 1 while the least similar have a weight of 0.

The net value $P_{i,w,m}$ is calculated as:

$$P_{i,w,m} = \sum_{k=1}^{n_{plots}} Y_{k|w,a} + B_{k,i|a} - O_i$$

where $Y_{k|w,a}$ is the estimated yield for plot *k*, using experiences from the farm(s) consulted for time path *w* and given land use choice *a*; $B_{k,i|a}$ is the sum of agglomeration and base payments for plot *k* in period *i* given the land use choice *a* and an expectation of the number of neighbors that would have adopted by year *i*; and O_i is the estimated sum of offers that would have been made to neighbors in order to encourage their adoption forward (see *Encouragement* submodel). For a given neighbor that has not yet adopted conservation practice, the expected offer out is drawn as a random value between 0 and the difference between the perceived value of their current portfolio, and the lowest-valued portfolio considered by the farm during their past turn that includes conservation practice. That is to say, in estimating their own expected costs for nudging, the farm assumes they would need to provide some fraction of the necessary amount to make conservation practice the most appealing option (Note – this is only the estimation of encouragement used for the farm's expected utility; the actual encouragements made are explained in the *Encouragement* submodel). For the first period in the time path, the net value would also include offers received by farm *j* conditional on adoption, and any costs associated with switching the land-use choice for plot *k*.

The algorithm for calculating expected utility is summarized as Loop 4 (Figure 5)

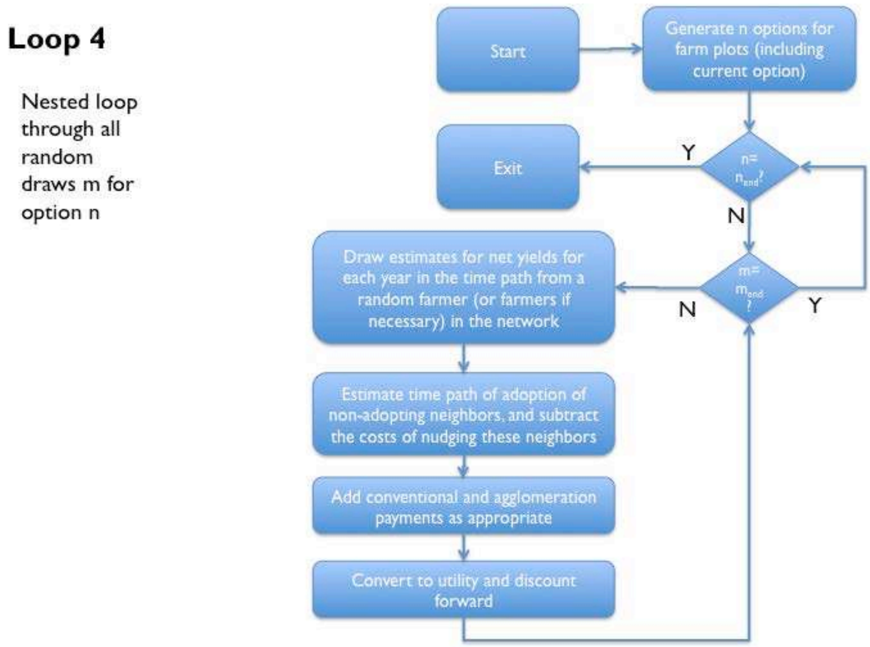

Figure 5 – Loop 4

Encouragements

An assumption in this model is that where farms stand to benefit from additional agglomeration payments when their neighbors also adopt conservation practice (or cheat), they provide an ‘encouragement’ in the form of an offered side payment to neighbors that have not yet adopted, conditional on their adoption of conservation practice. Some farms will honor the offer if the neighbor cheats rather than adopts; this is a farm-specific parameter drawn at initialization and fixed for the duration of the simulation.

Farms make or update their offers to non-adopting neighbors at the end of every interaction loop in which their intent for the upcoming season includes conservation practice or cheating (Loop 2). The maximum amount $O_{max,j}$ a farm j will offer to another farm is farm-specific, and is some fraction of a single-season’s agglomeration payment value between 0 and 1. The first offer made by a farm j to a new, non-adopting neighboring farm is a random fraction of $O_{max,j}$; subsequent offers made to the same farm update the offer by an additional, randomly drawn fraction, up to the maximum of O_{max}.

When farms adopt conservation practice, or when they cheat (in the case that the farm offering them a side payment will honor the offer under cheating), they receive outstanding offered side payments from those farms that had offered them.

Rainfall

Rainfall is treated very simply in this model – the same static normal distribution is used to draw rainfall for each plot in each year. As rainfall is plot-specific, there is room in model extensions to explore different spatio-temporal relationships for rainfall, but in the current model it is kept very simple.

Table 1 – Modeling parameters used in experiments

Modeling Parameters - Varied in Monte Carlo Analysis	Minimum	Maximum	Units
Number of interaction steps per decision time step	5	9	/year
Rainfall Mean	250	500	mm/year*
Rainfall SD	50	100	mm/year*
Number of time steps / years considered in estimating expected utility	15	25	year
Transaction cost, switch to conservation practice	50	300	dollar*
Transaction cost, switch to conventional practice	50	300	dollar*
Average link strength for network	0.3	0.8	-
Plot Radius Mean	50	80	m*
Plot Radius SD	20	80	m*
Network link Kernel Radius	500	2000	m*
Plot Soil Quality SD	0.05	0.1	-
Farm Discount Rate Mean	0.02	0.1	-
Farm Discount Rate SD	0.01	0.05	-
Farm CRRA Mean	0.5	1.25	-
Farm CRRA SD	0.1	0.3	-
Farm Participation in Interaction Step Mean		0.9	-
Farm Participation in Interaction Step SD	0.1	0.2	-
Number of random draws to calculate expected utility	10	20	-
Perceived likelihood of a given neighbor adopting, for use when estimating expected utility	0.3	0.7	-
Number of plots per farm Mean	2	4	-
Number of plots per farm SD	0	3	-
Number of options considered Mean	6	14	-
Number of options considered SD	1	3	-
Probability of leaving landscape	0.01	0.04	-
Farm technical efficiency SD	0.05	0.1	-
Probability that a farm honors an offer to a neighbor that cheats	0	1	-

Modeling Parameters - Fixed in all simulations	Value	Units
Fraction of landscape filled	0.6	-
Landscape X	3000	m*
Landscape Y	3000	m*
Years in simulation	30	year
Number of sponsored pilot study members	10	farm
Start year for early adoption pilot	5	-
End year for early adoption pilot	8	-
Length of subsidy	6	years

Modeling Parameters - Varied in Policy Sweep	Minimum	Increment	Maximum	Units
Base Payment	0	20	100	dollar/year*
Agglomeration Payment	0	20	100	dollar/year*
Penalty for Cheating	1000	500	2000	dollar/year*
Random monitoring Probability	0.2	0.2	0.8	-

Sensitivity Analysis

This sensitivity analysis across social and environmental variables in our model was used as a basis for developing the Monte Carlo Analysis used in the current study.

As part of a team discussion, study authors identified 26 of the 36 tested variables as appearing (visually) to have an impact on the outcomes of interest (area under adoption, and program spending on payments). The ranges for these 26 variables used in the Monte Carlo Analysis are detailed in Table 1 in the main paper.

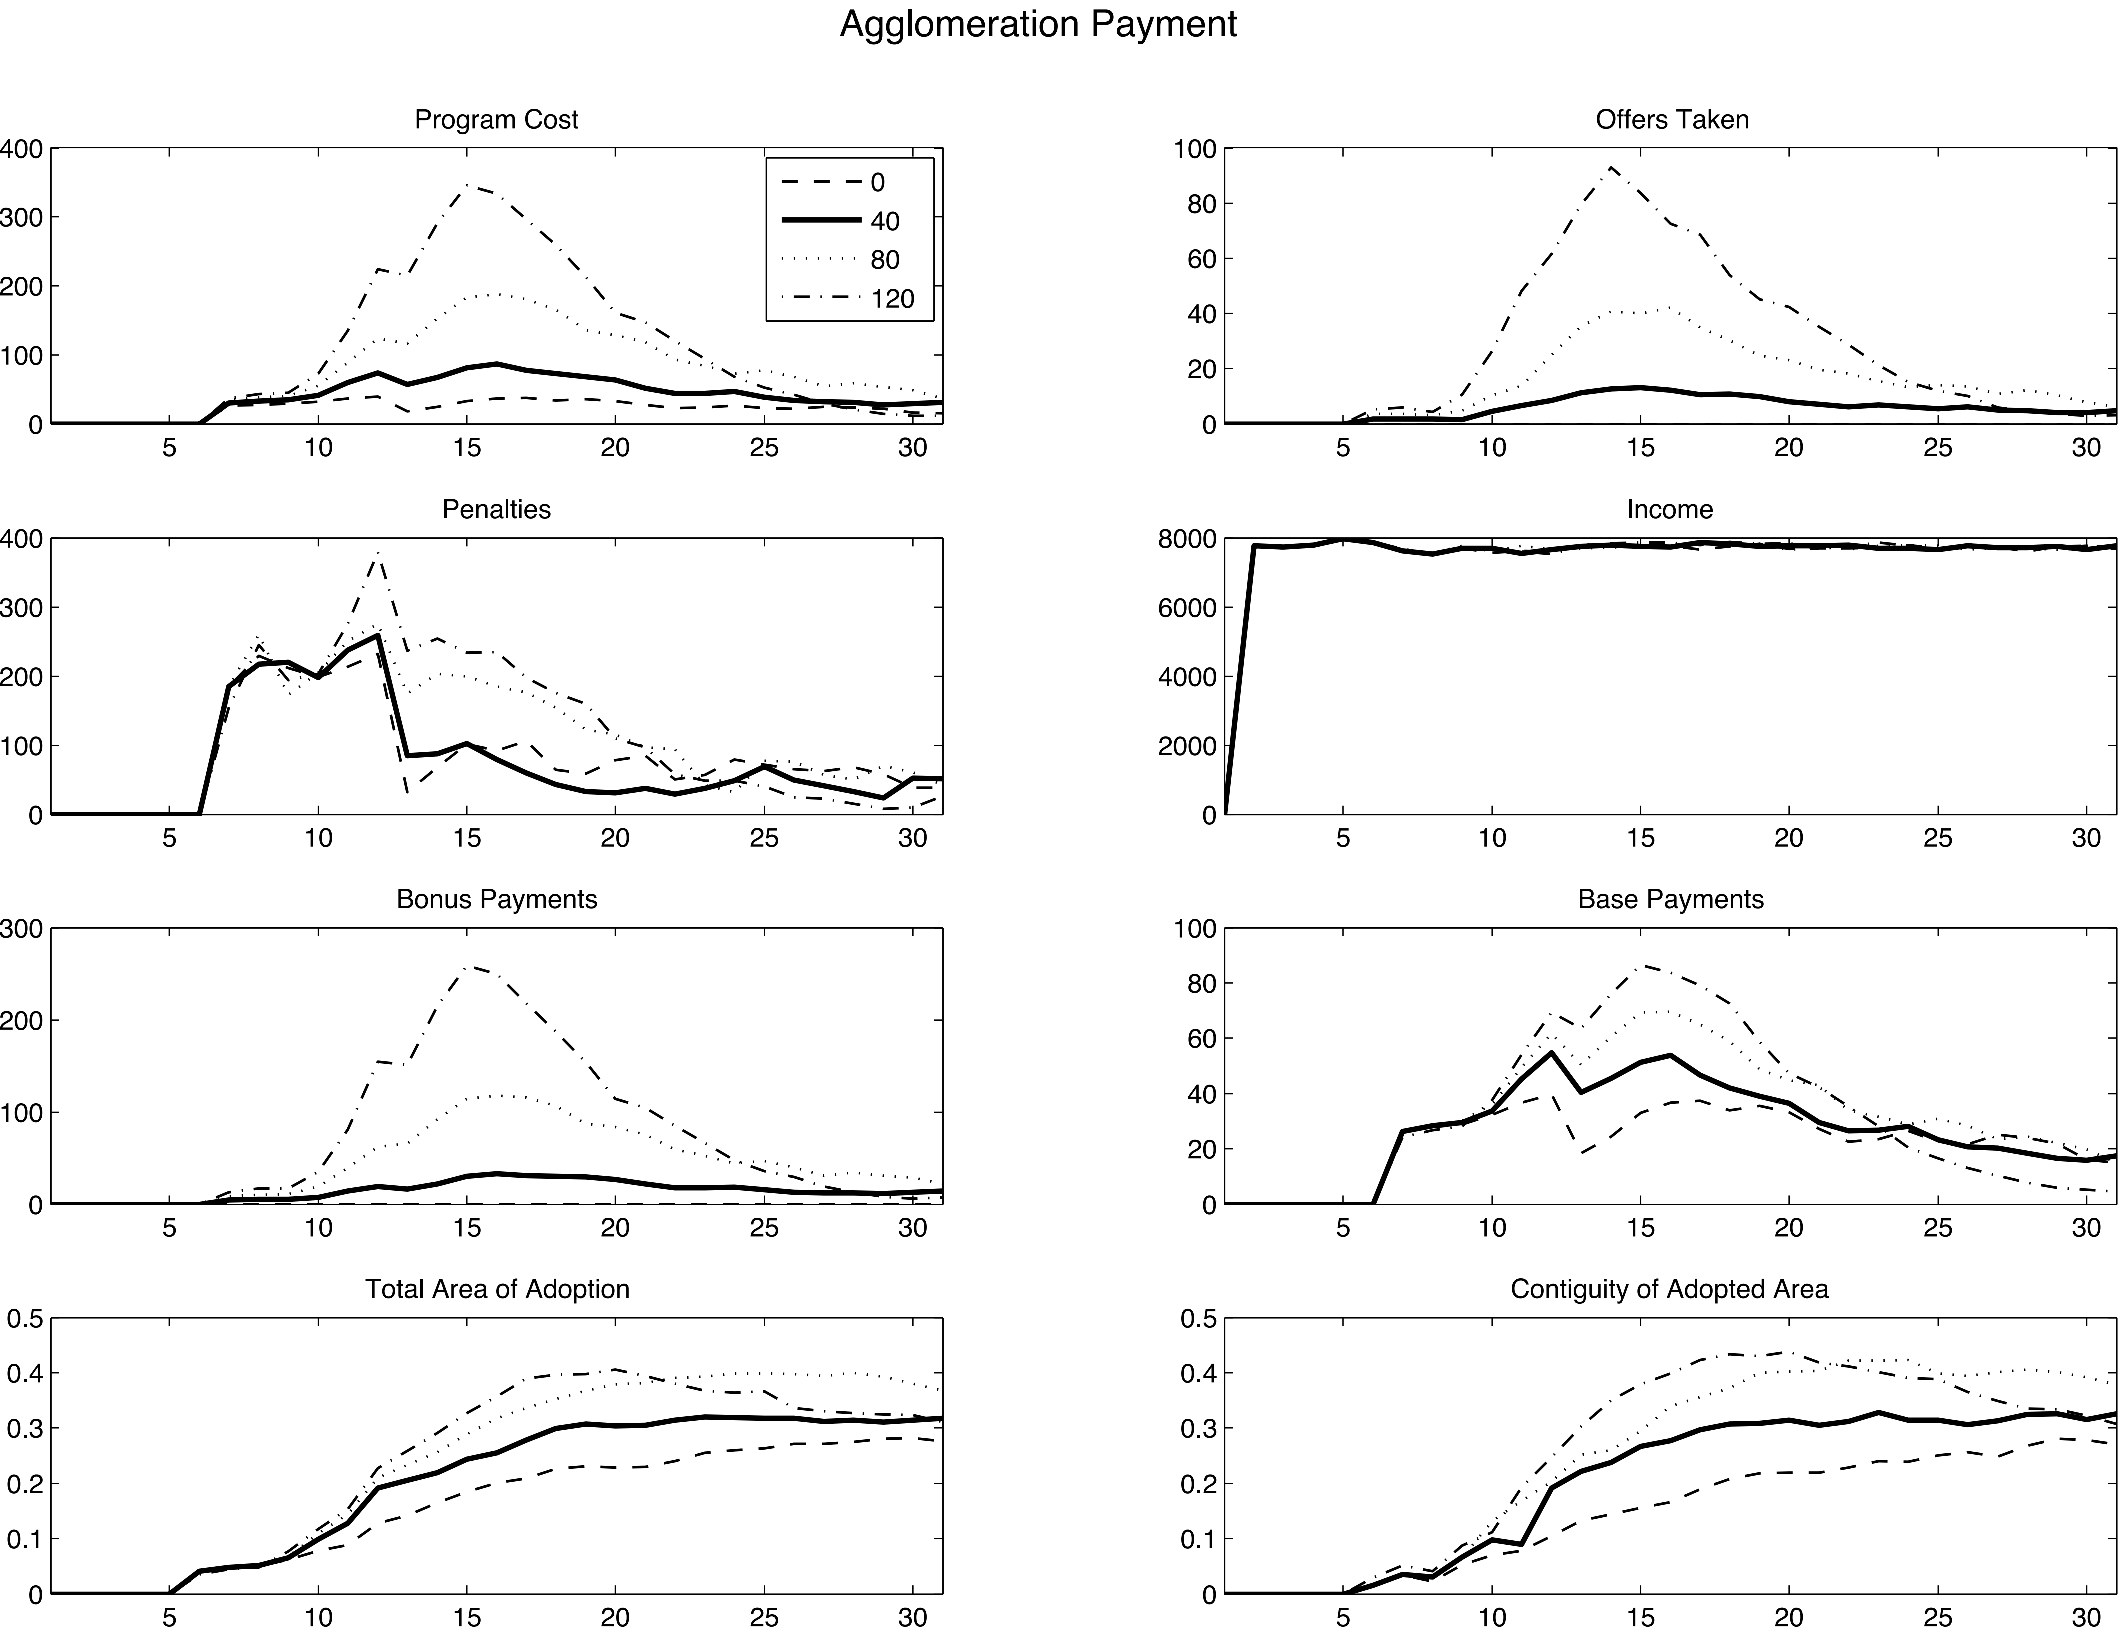
Agglomeration Payment
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
0
40
80
120

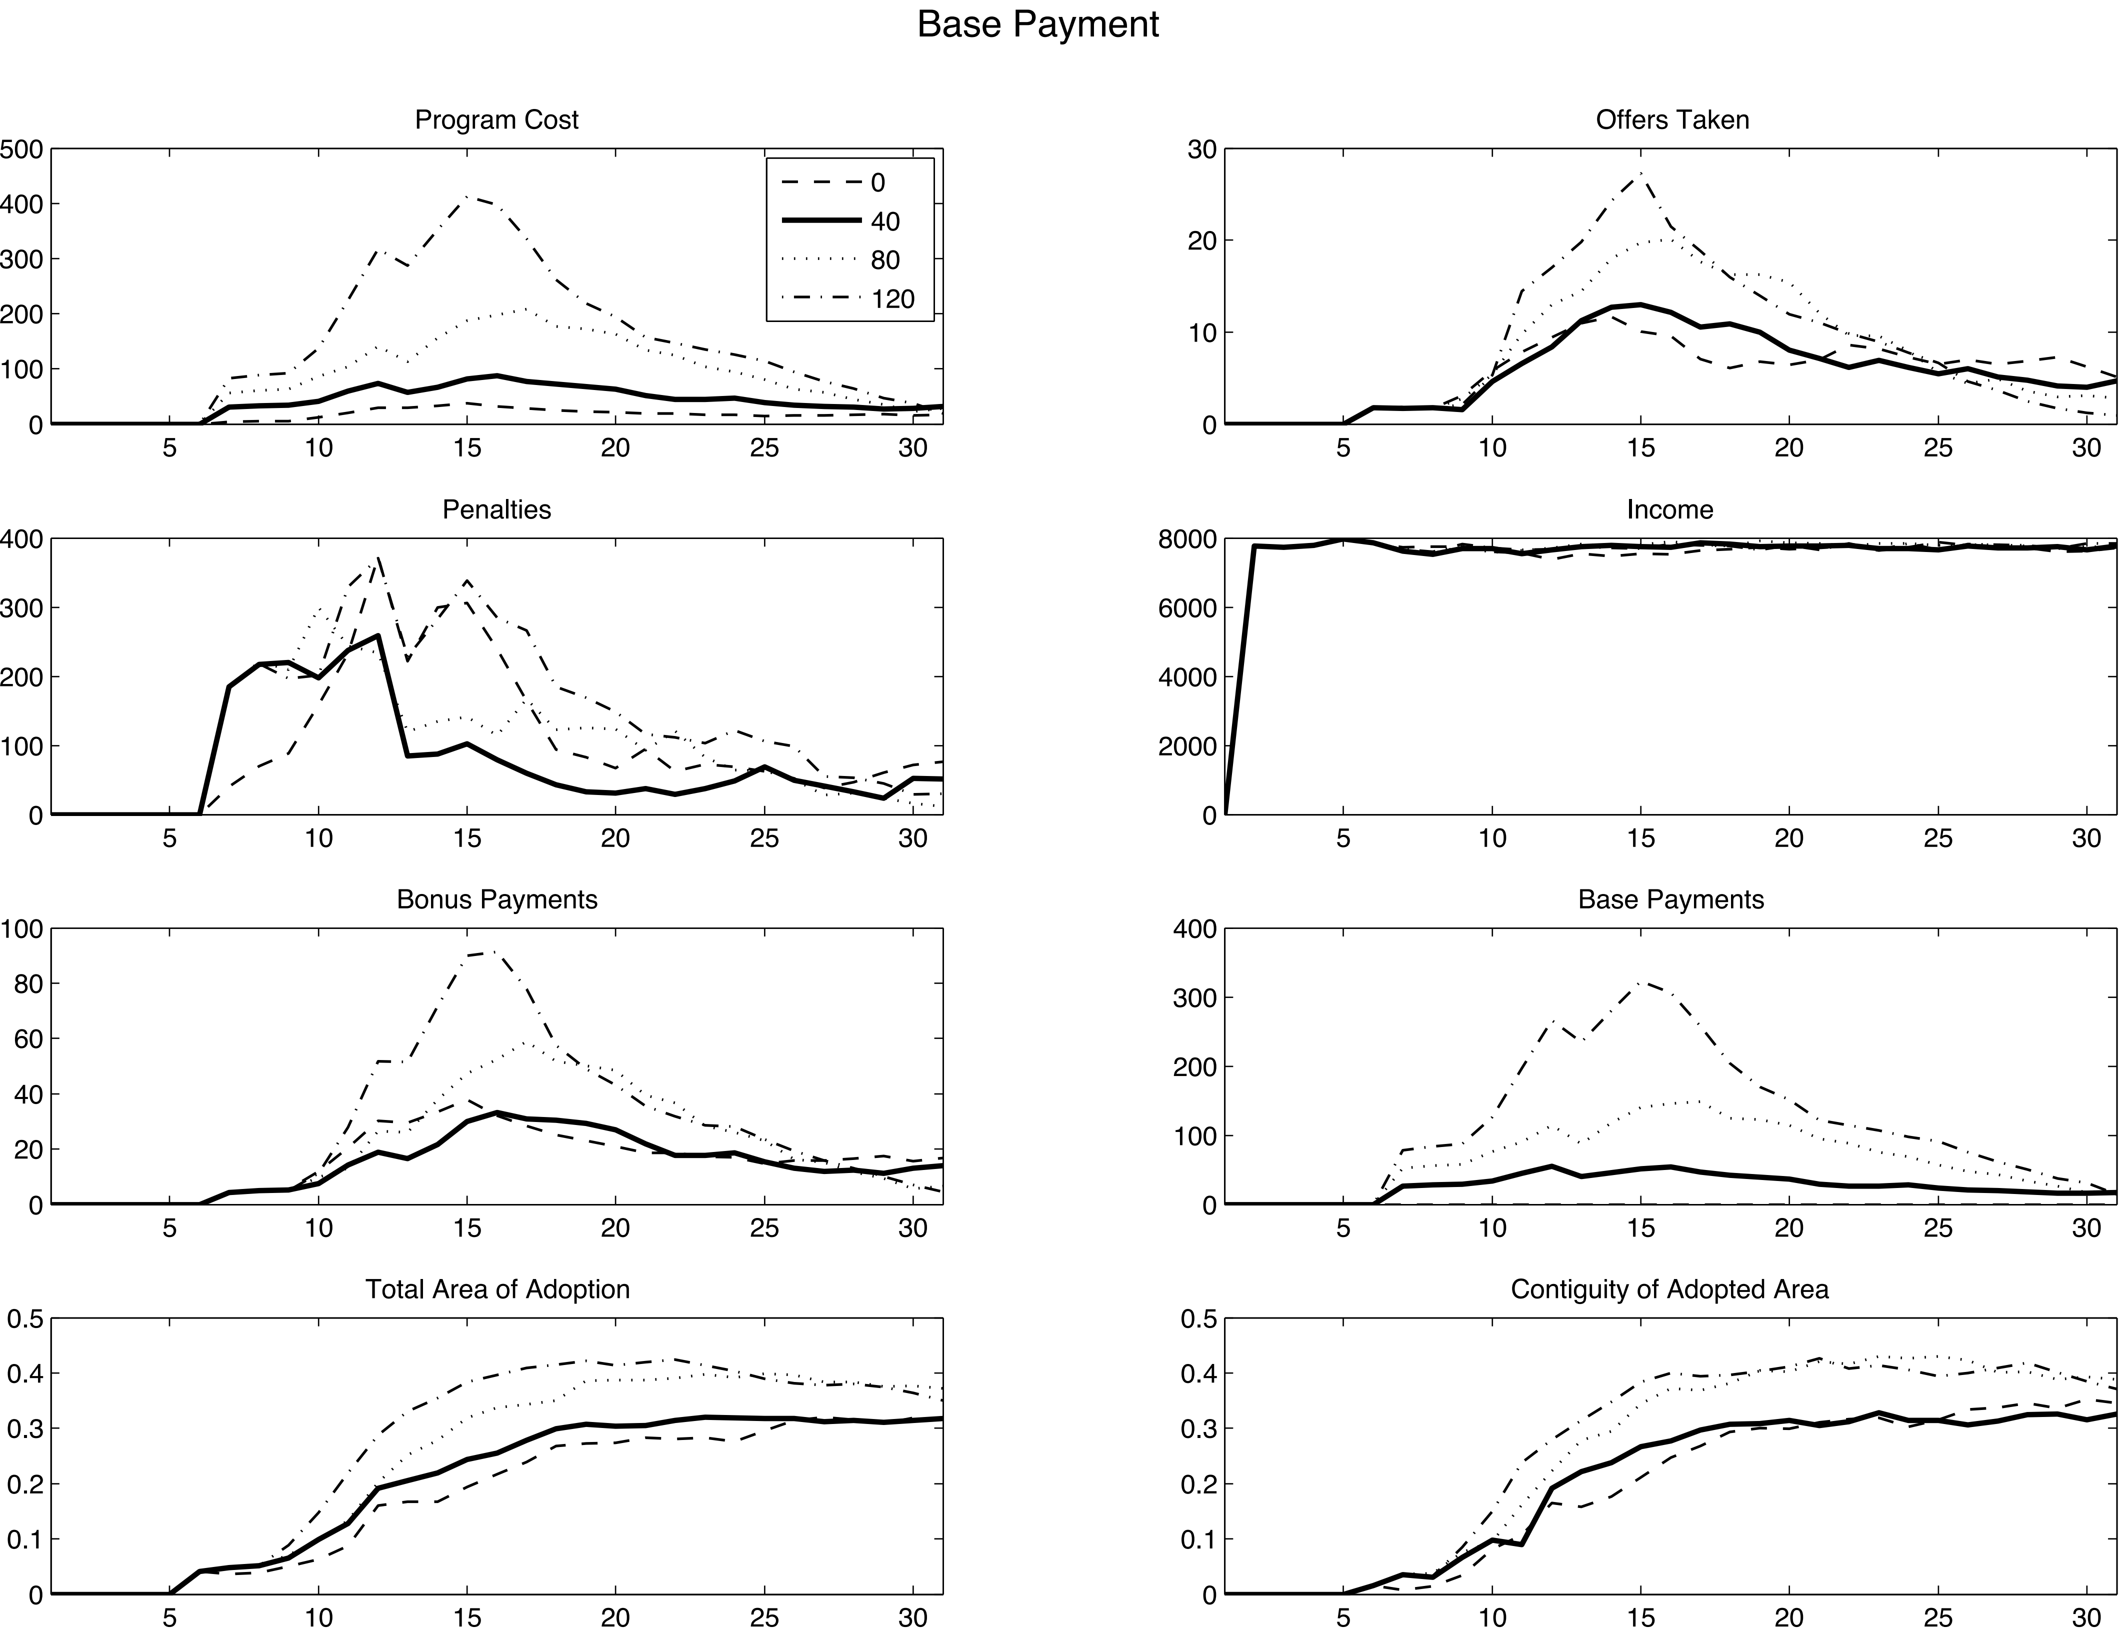

Base Payment
Program Cost
0
40
80
120
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

Chance Adopt

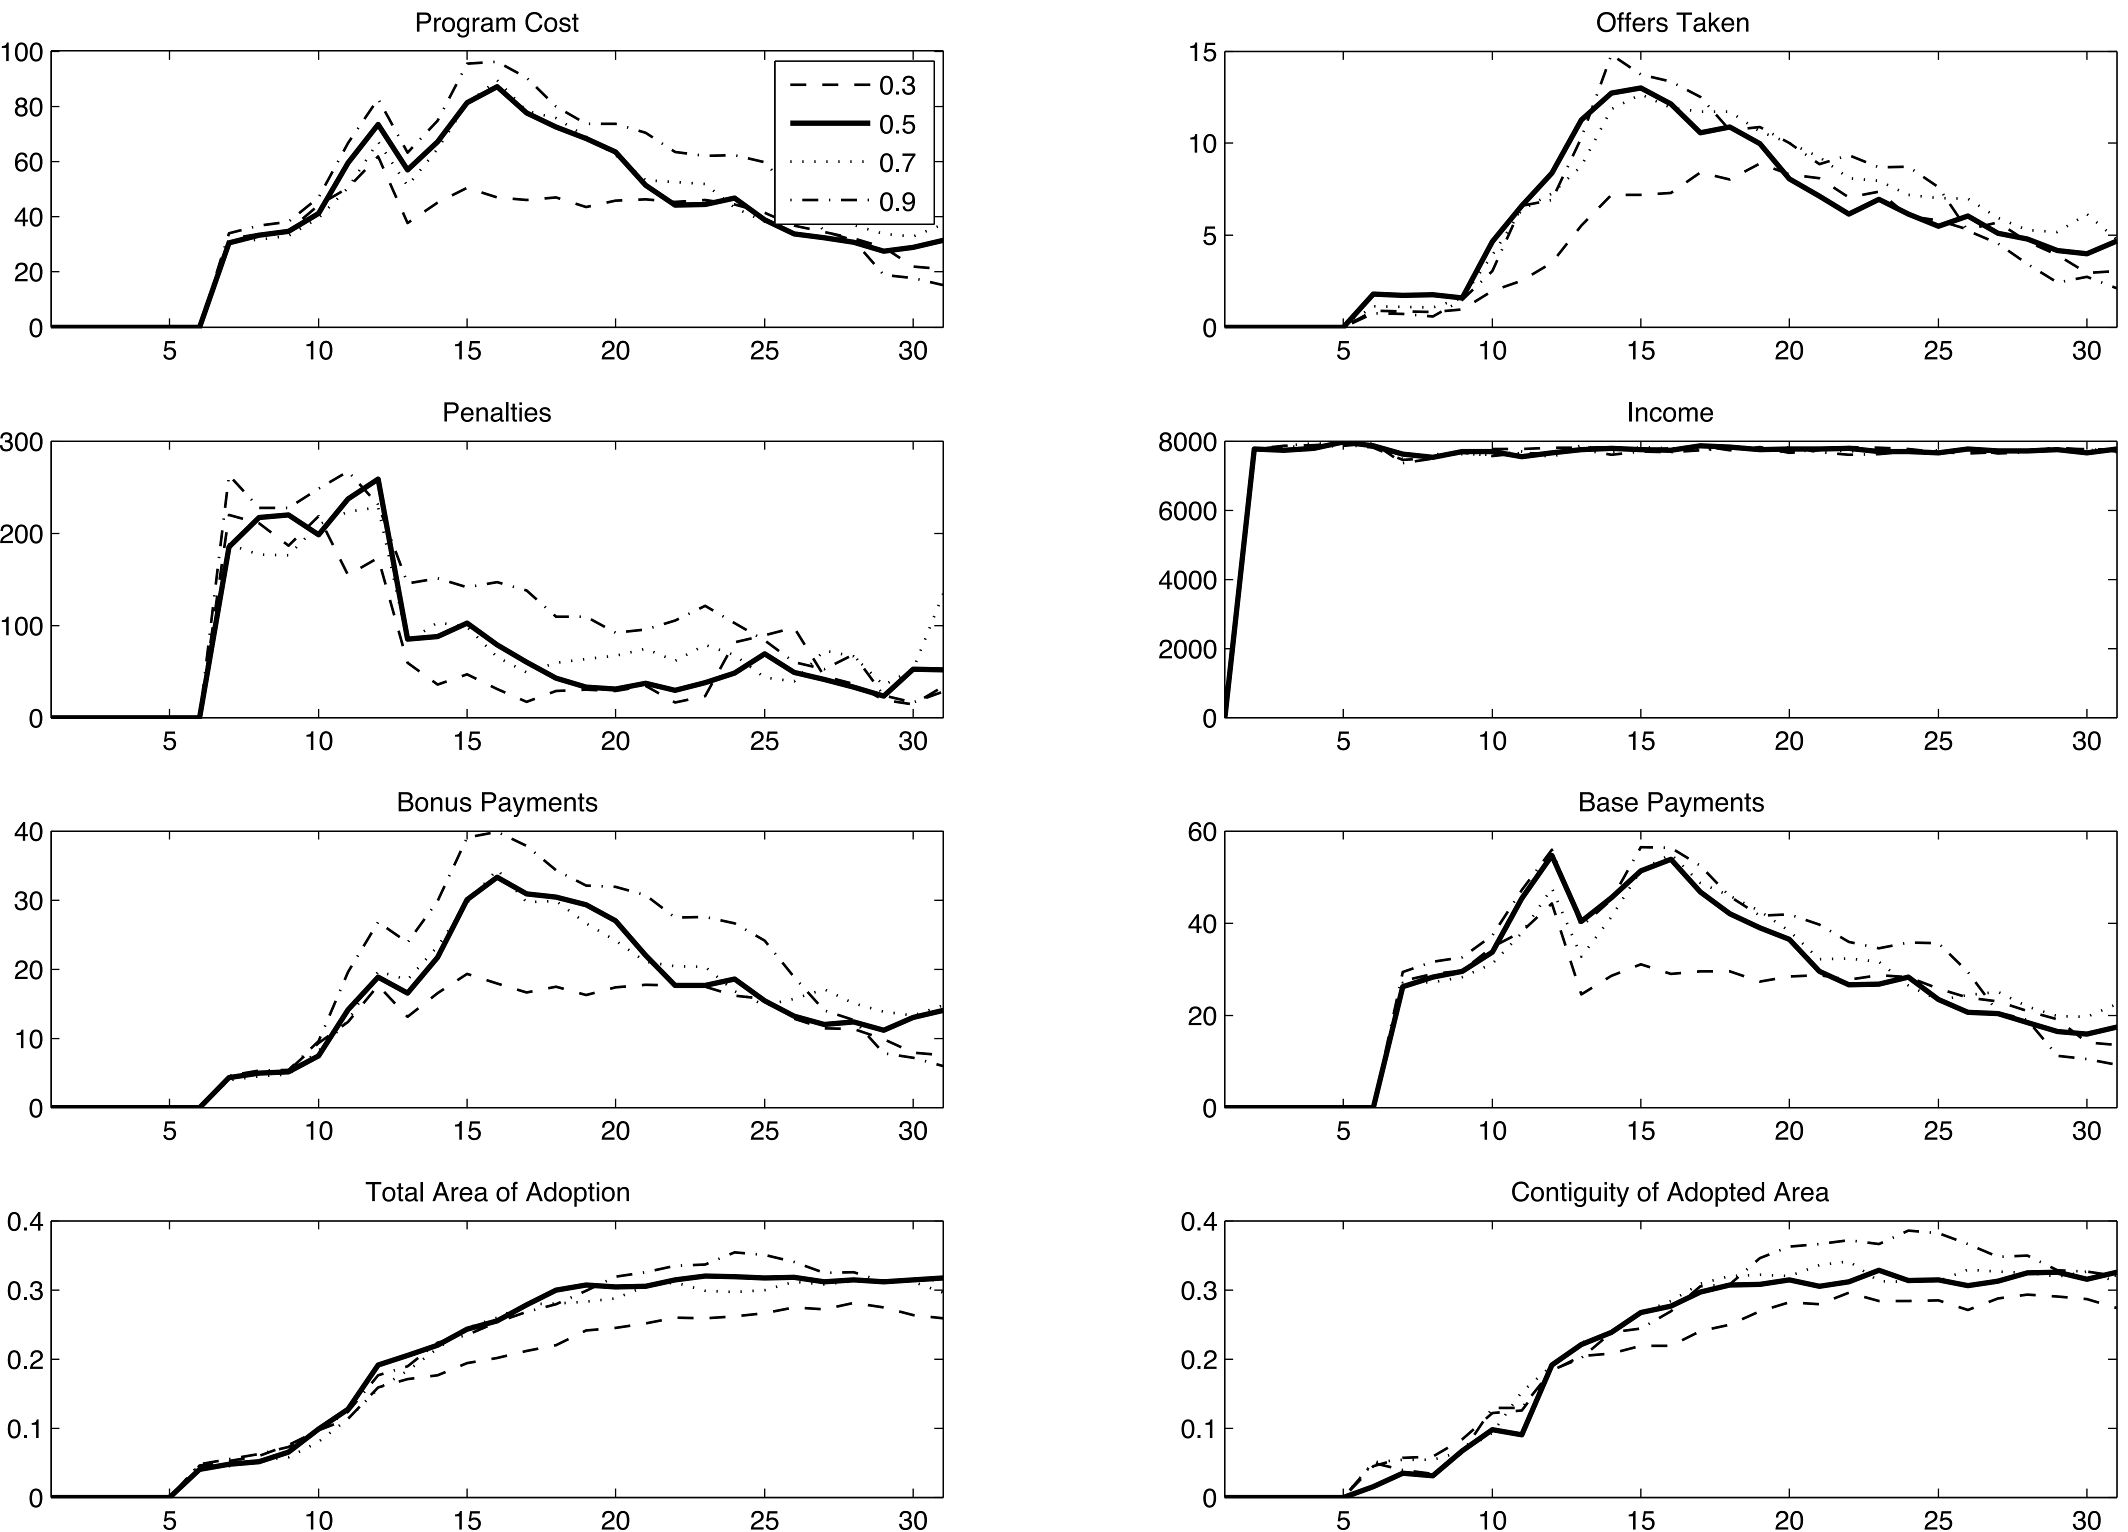

Climate Mean

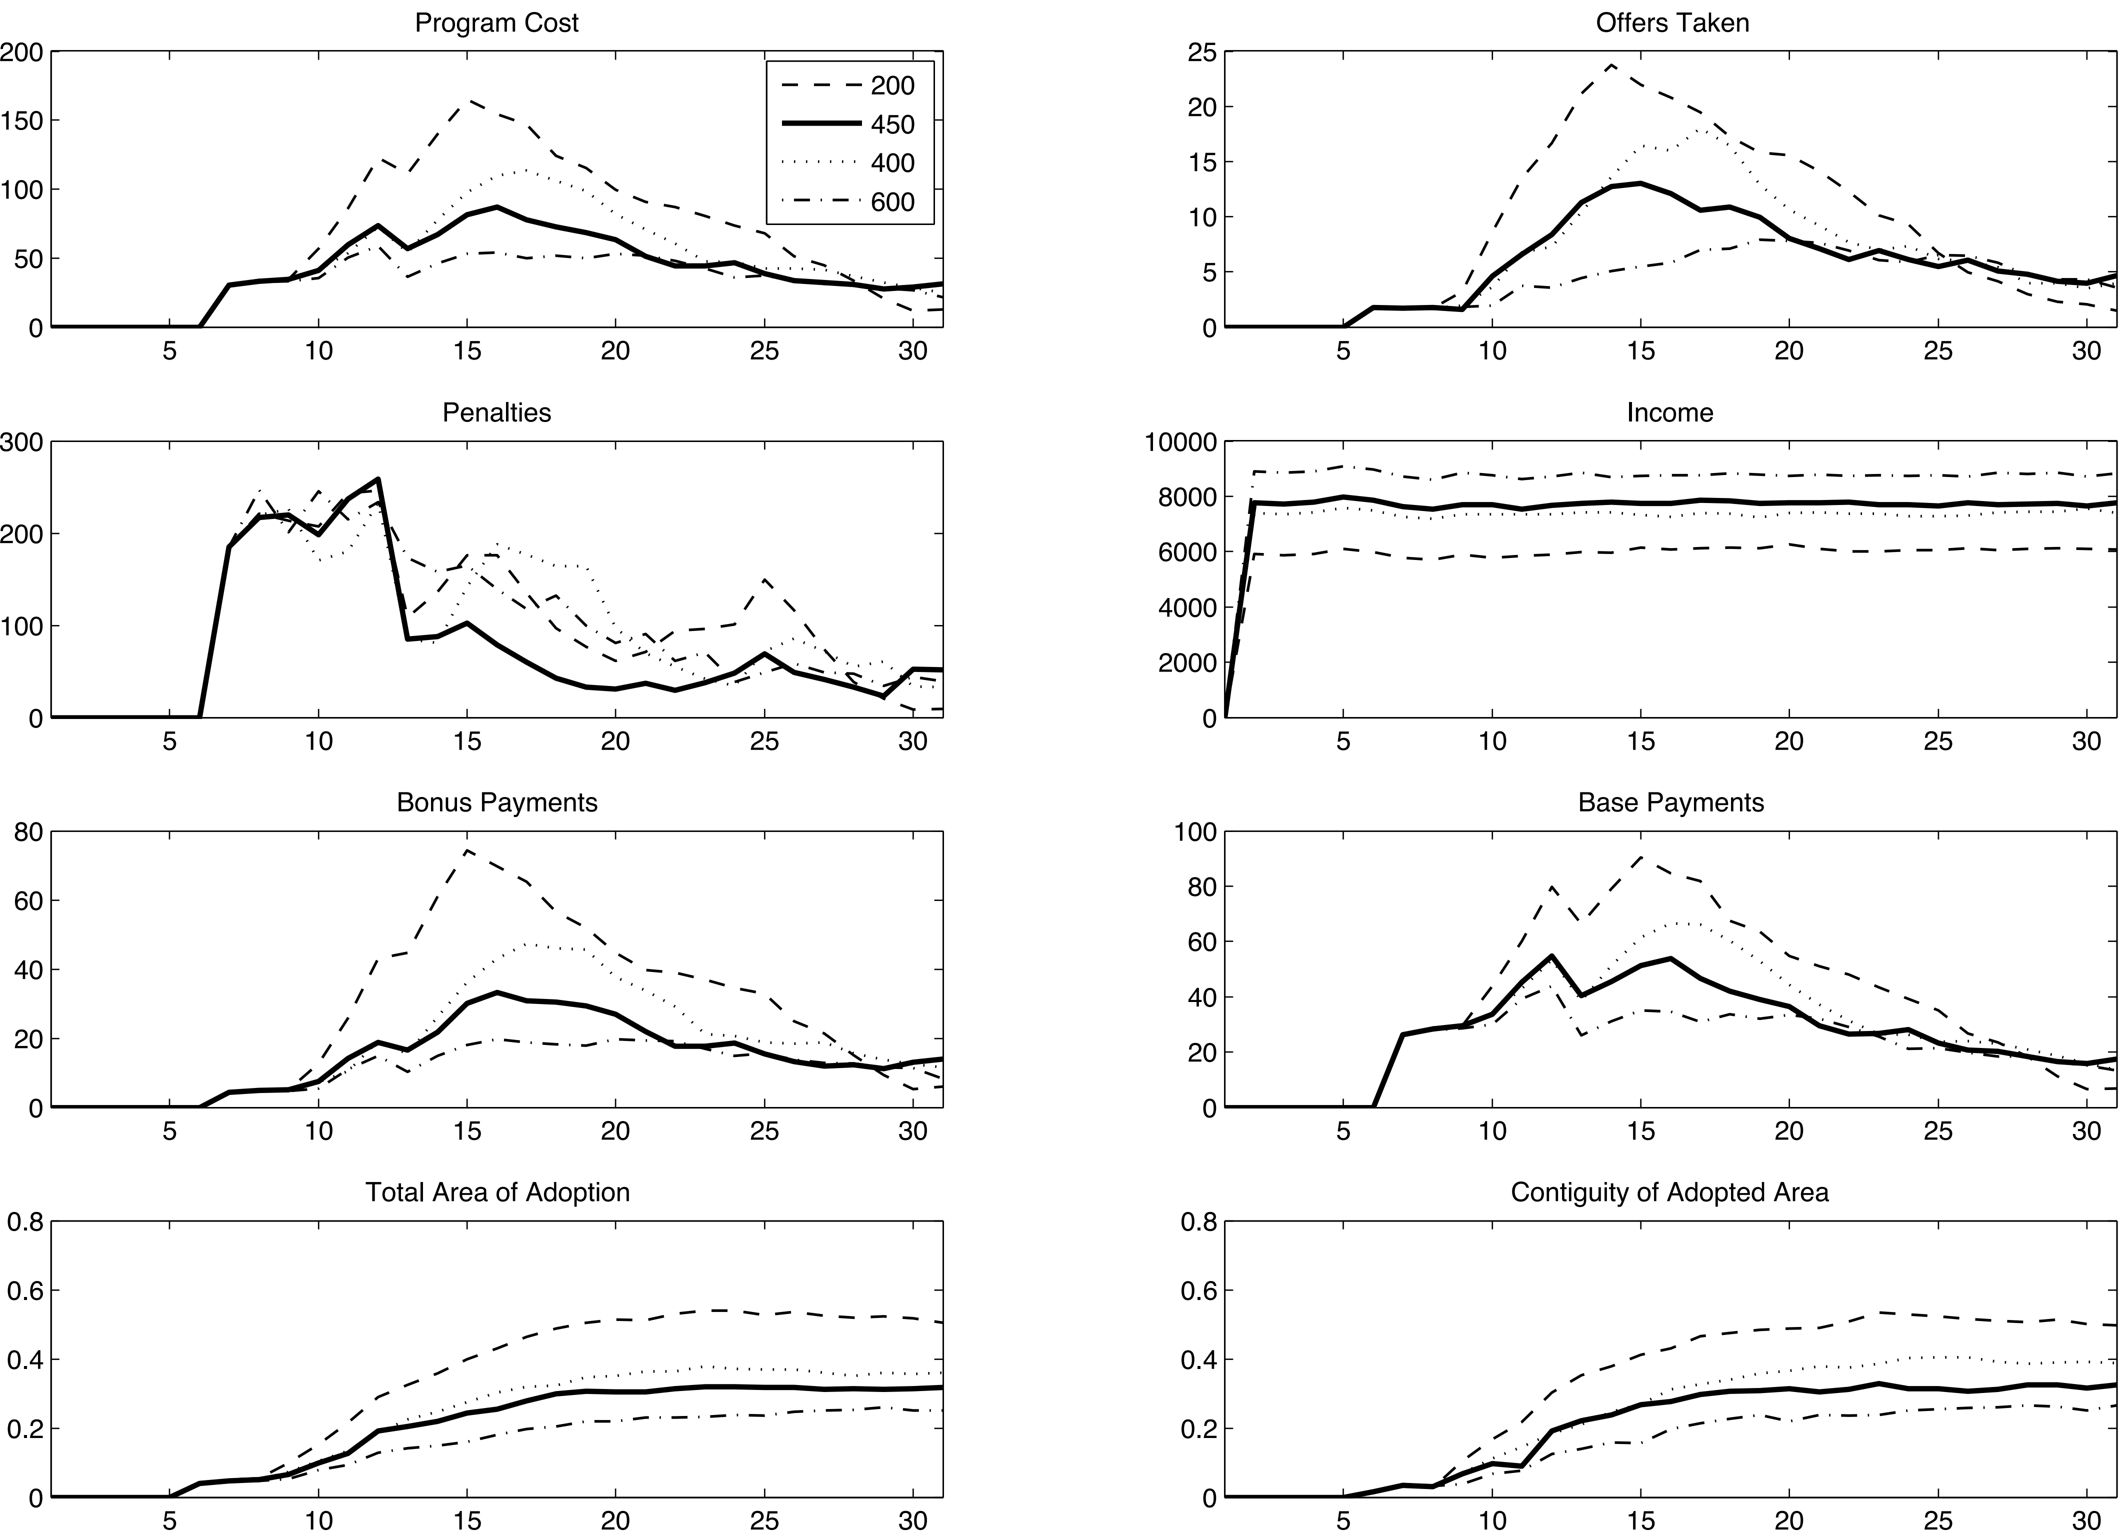

Climate SD

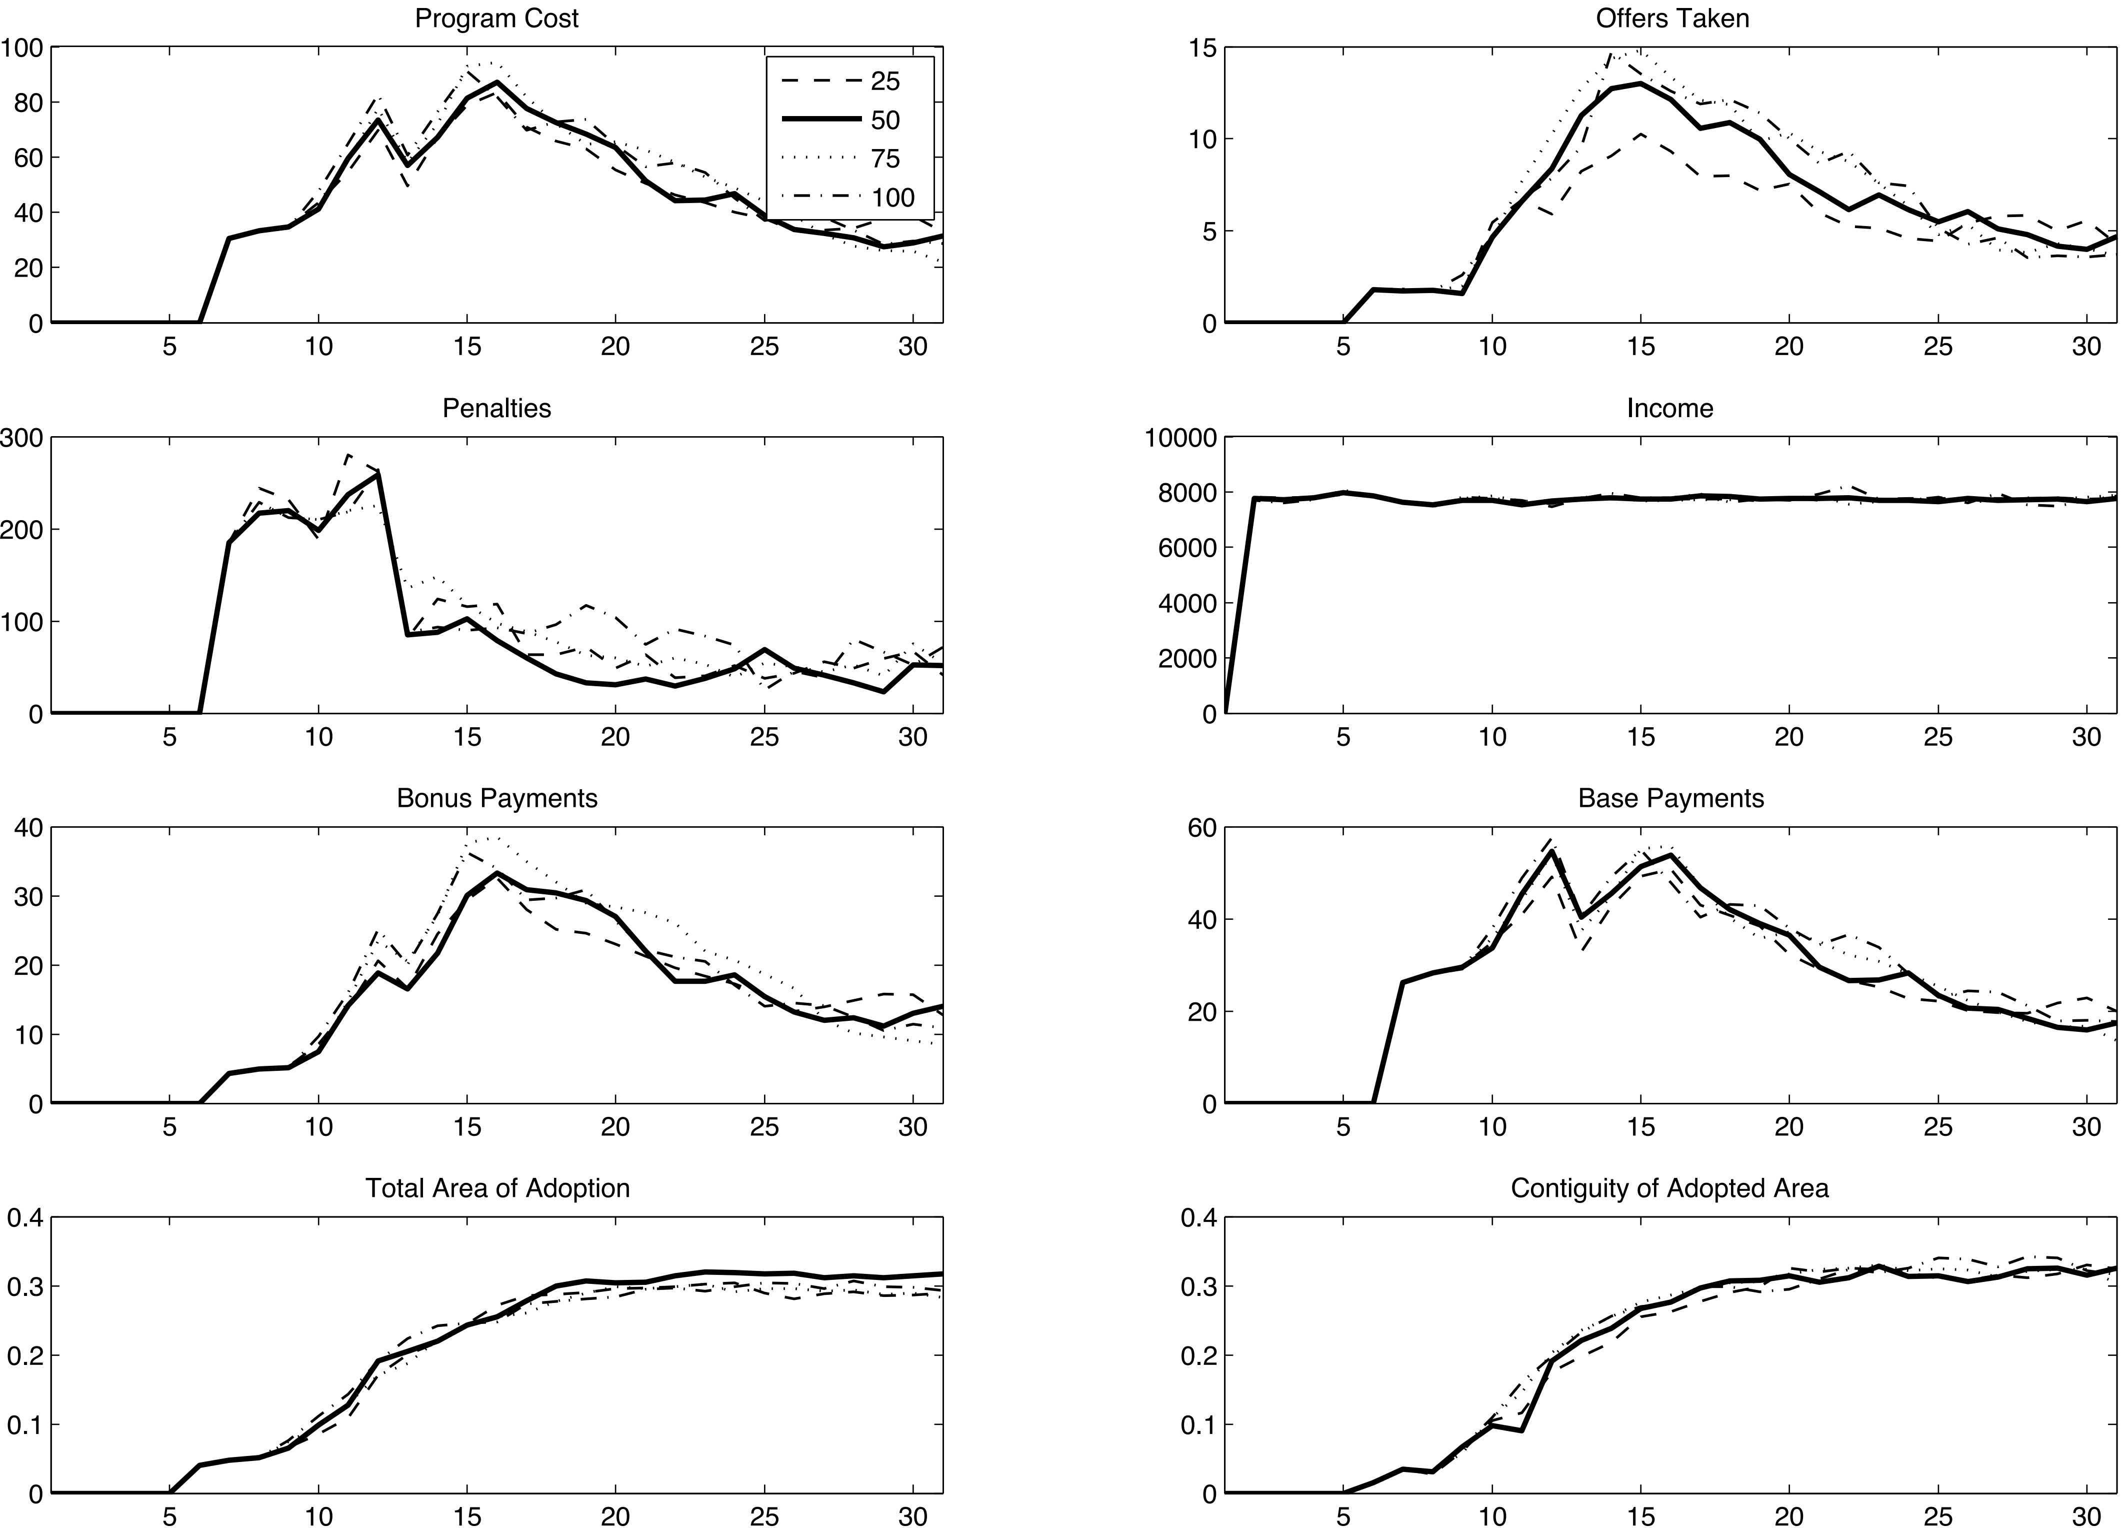

Program Cost
25
50
75
100
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

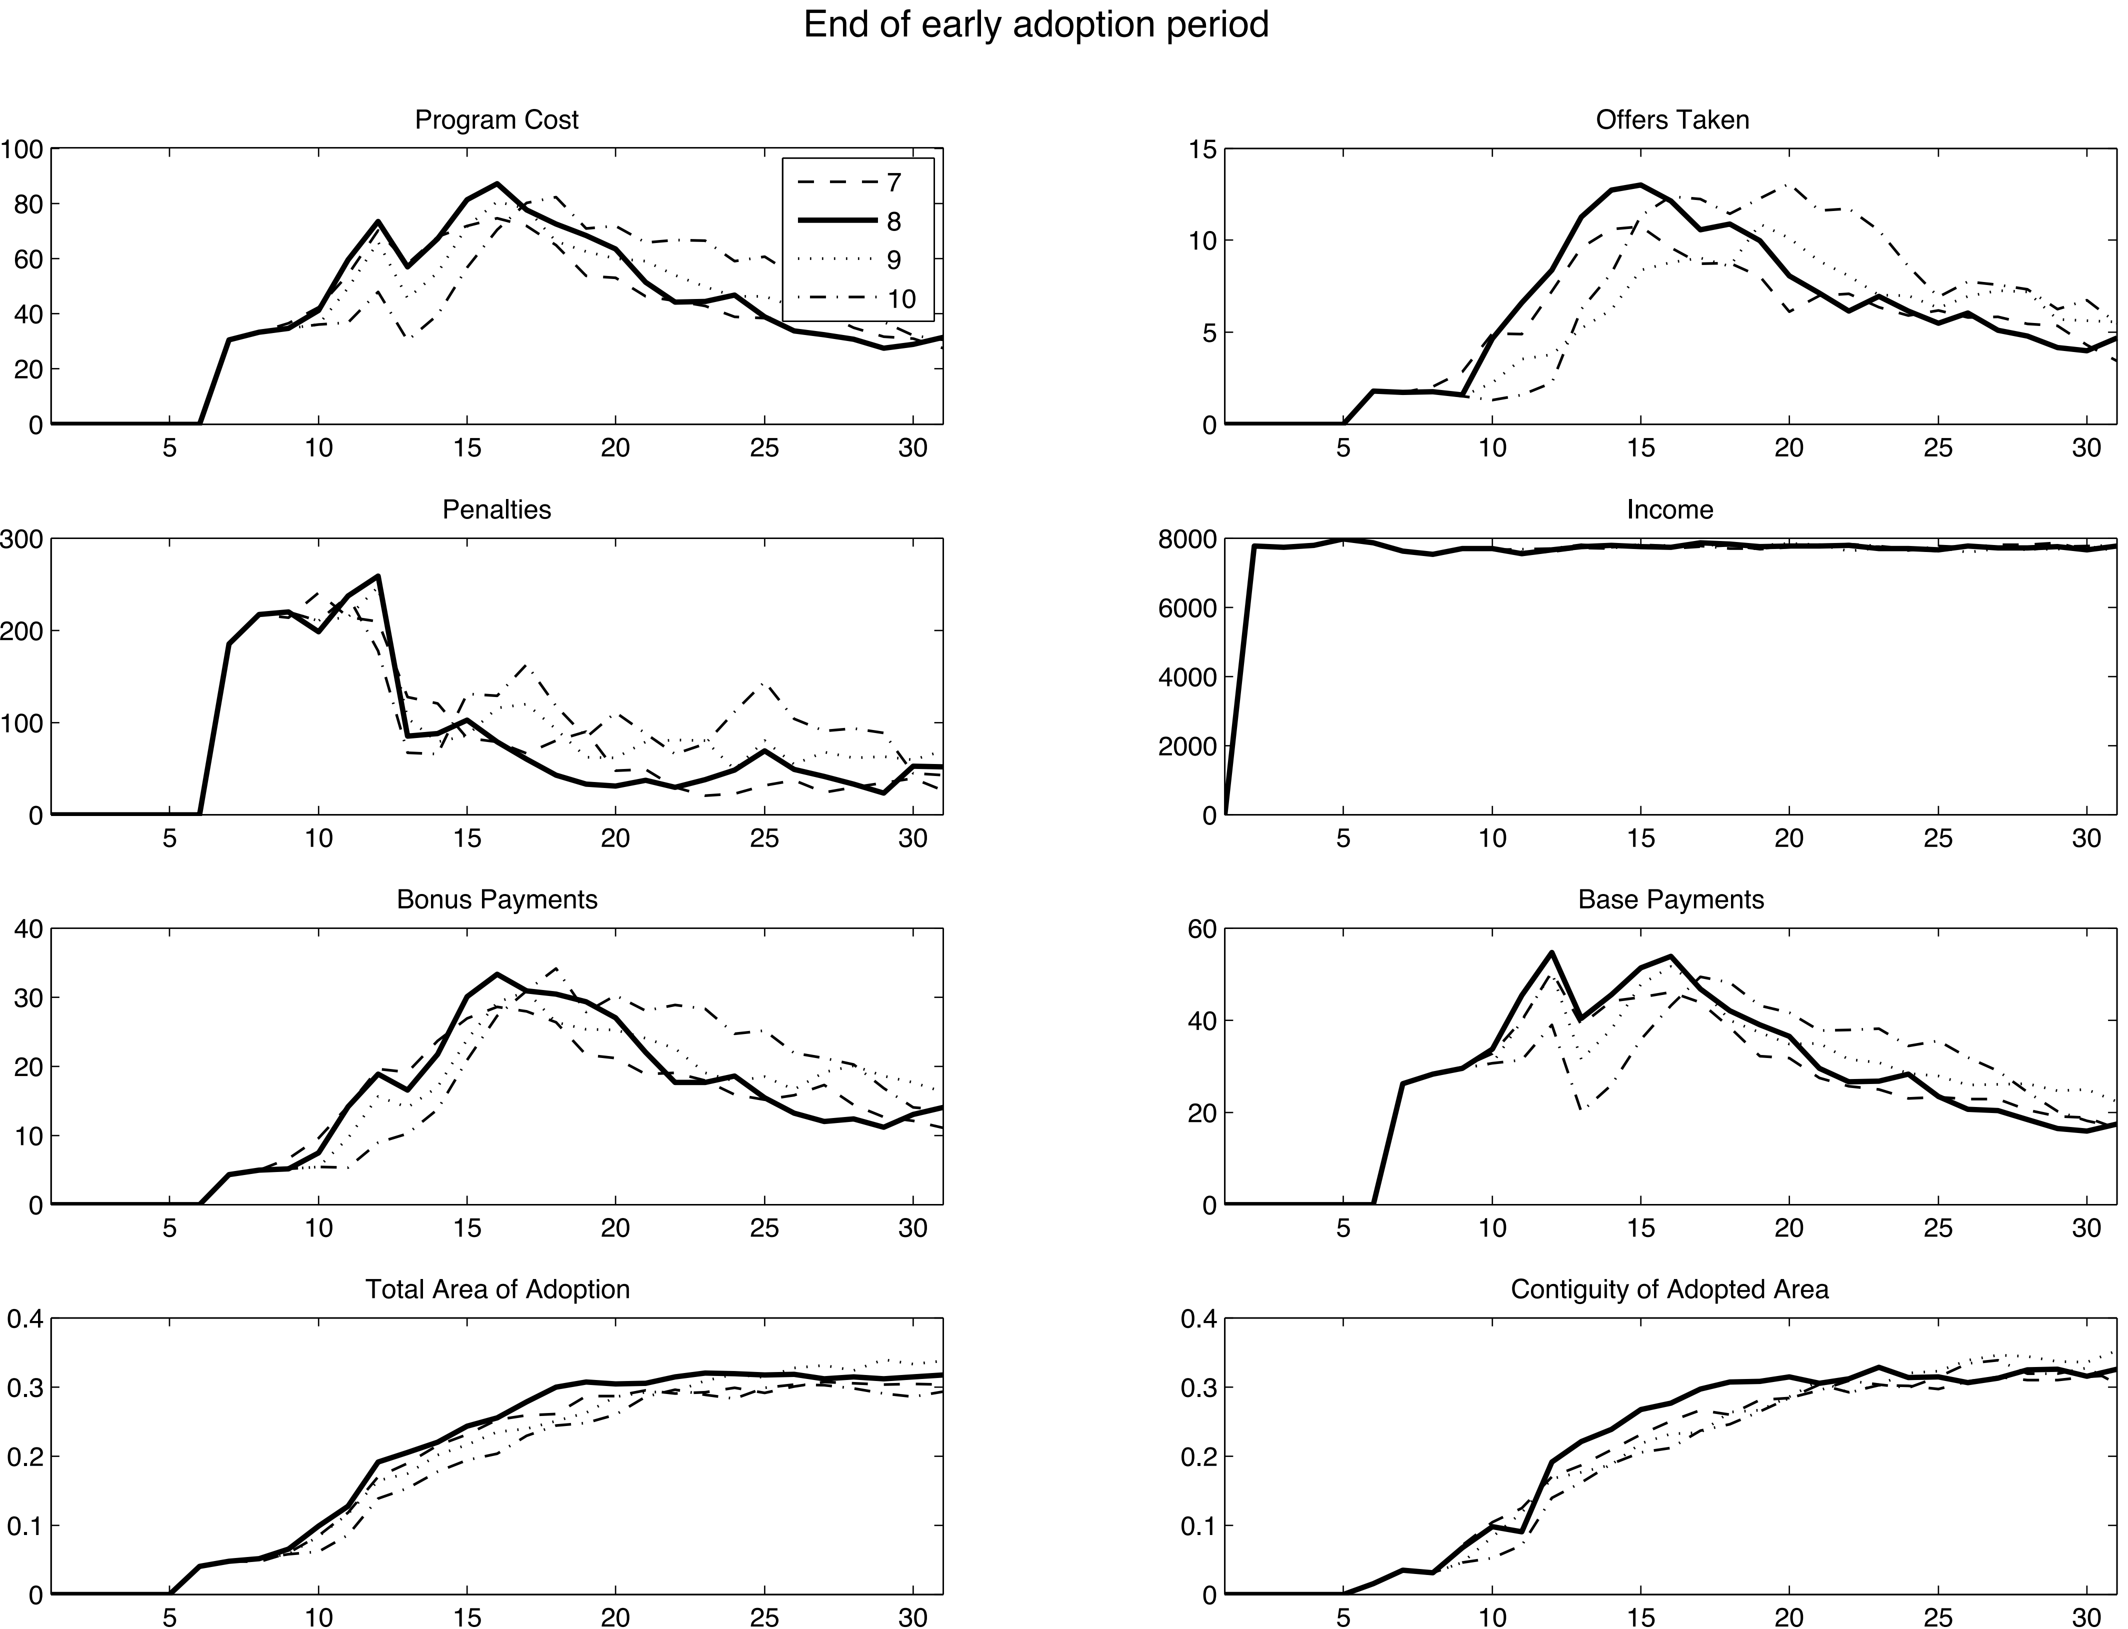
End of early adoption period
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
7
8
9
10

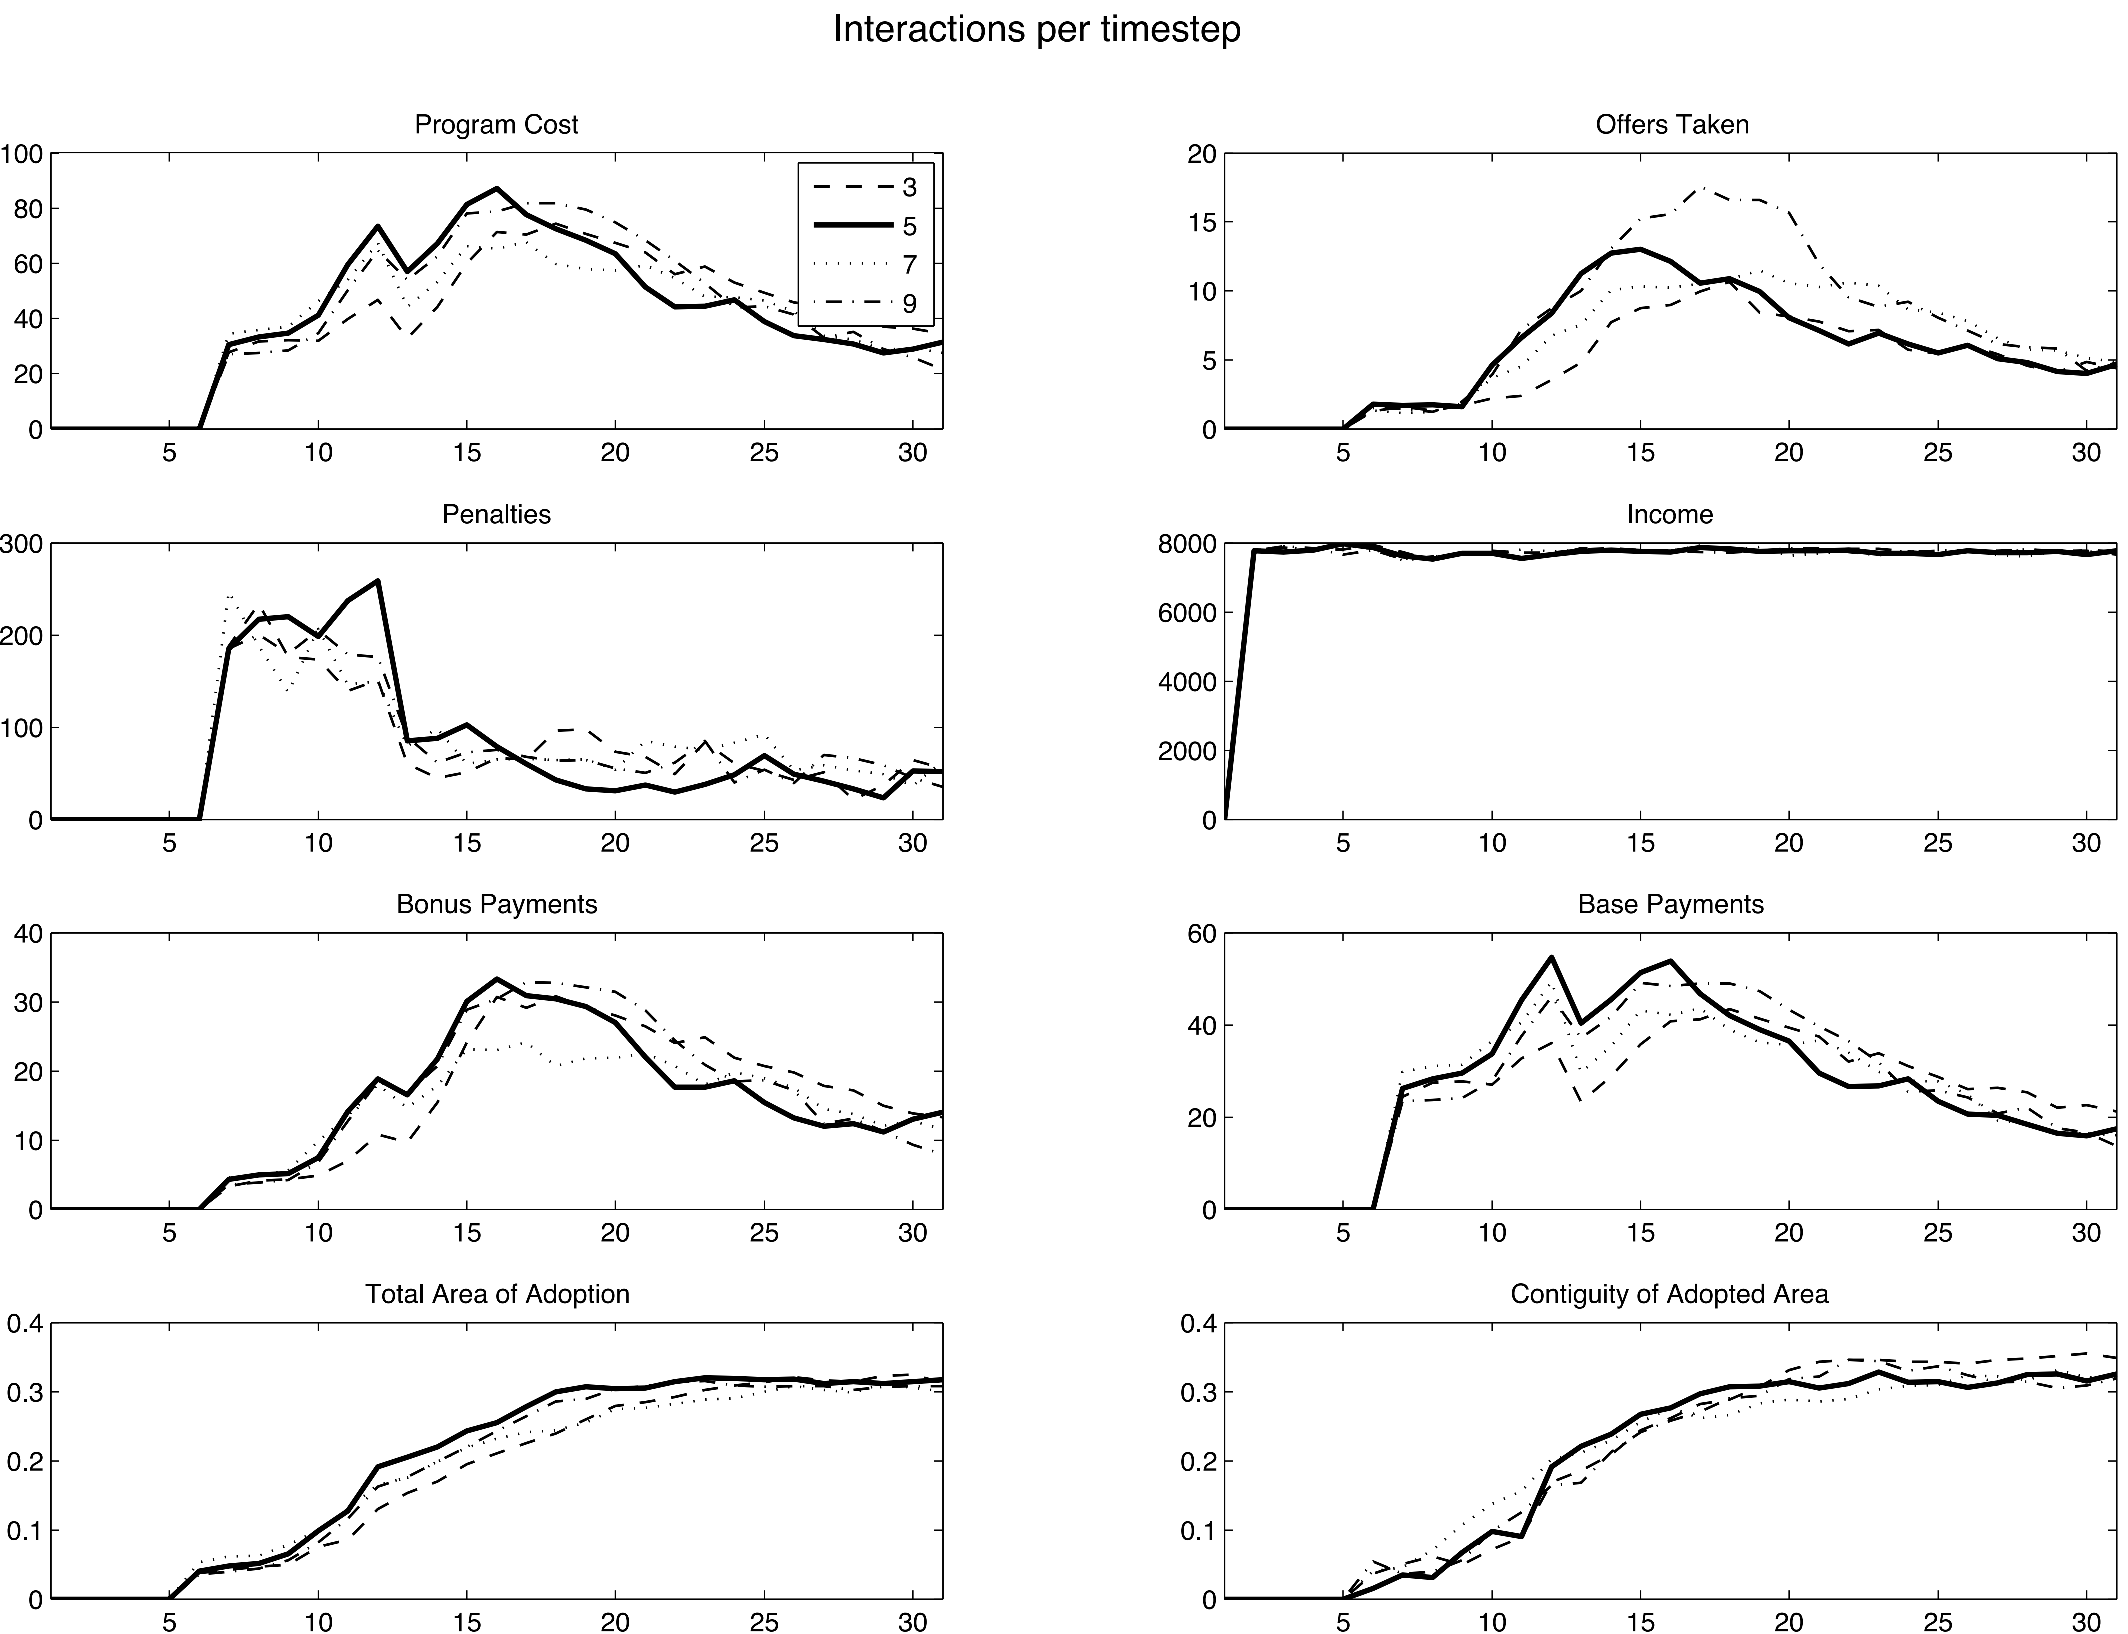
Interactions per timestep
Program Cost
3
5
7
9
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

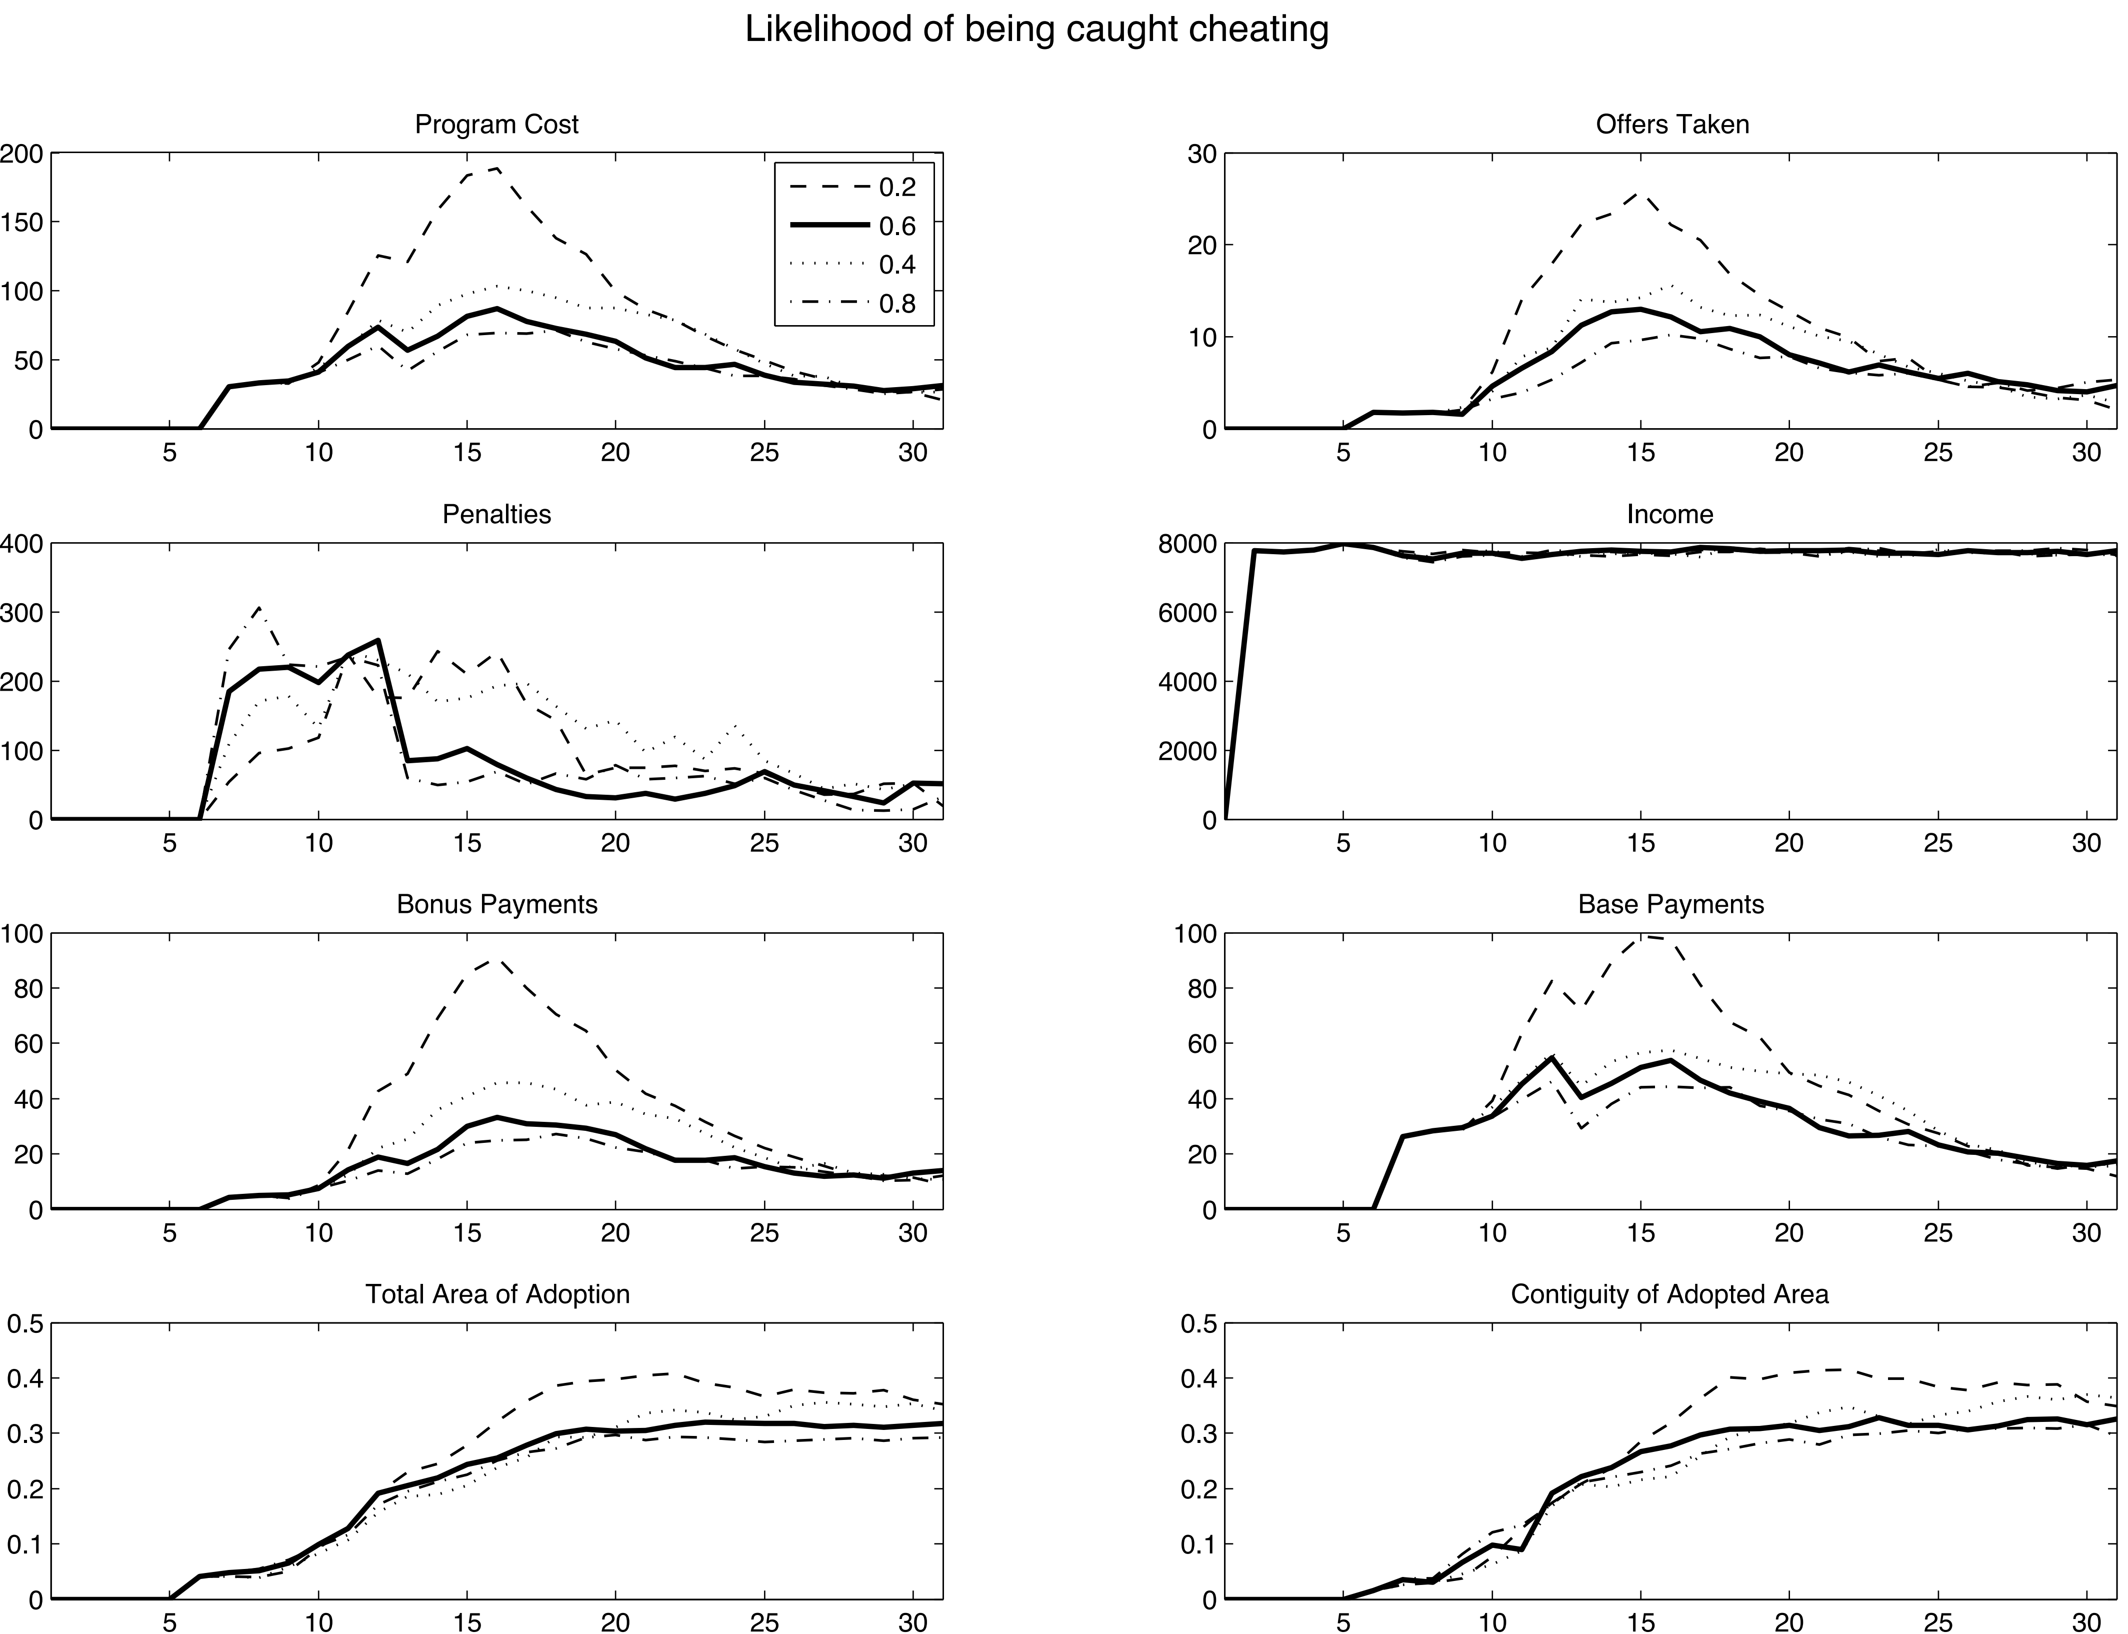

Likelihood of being caught cheating
Program Cost
0.2
0.6
0.4
0.8
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

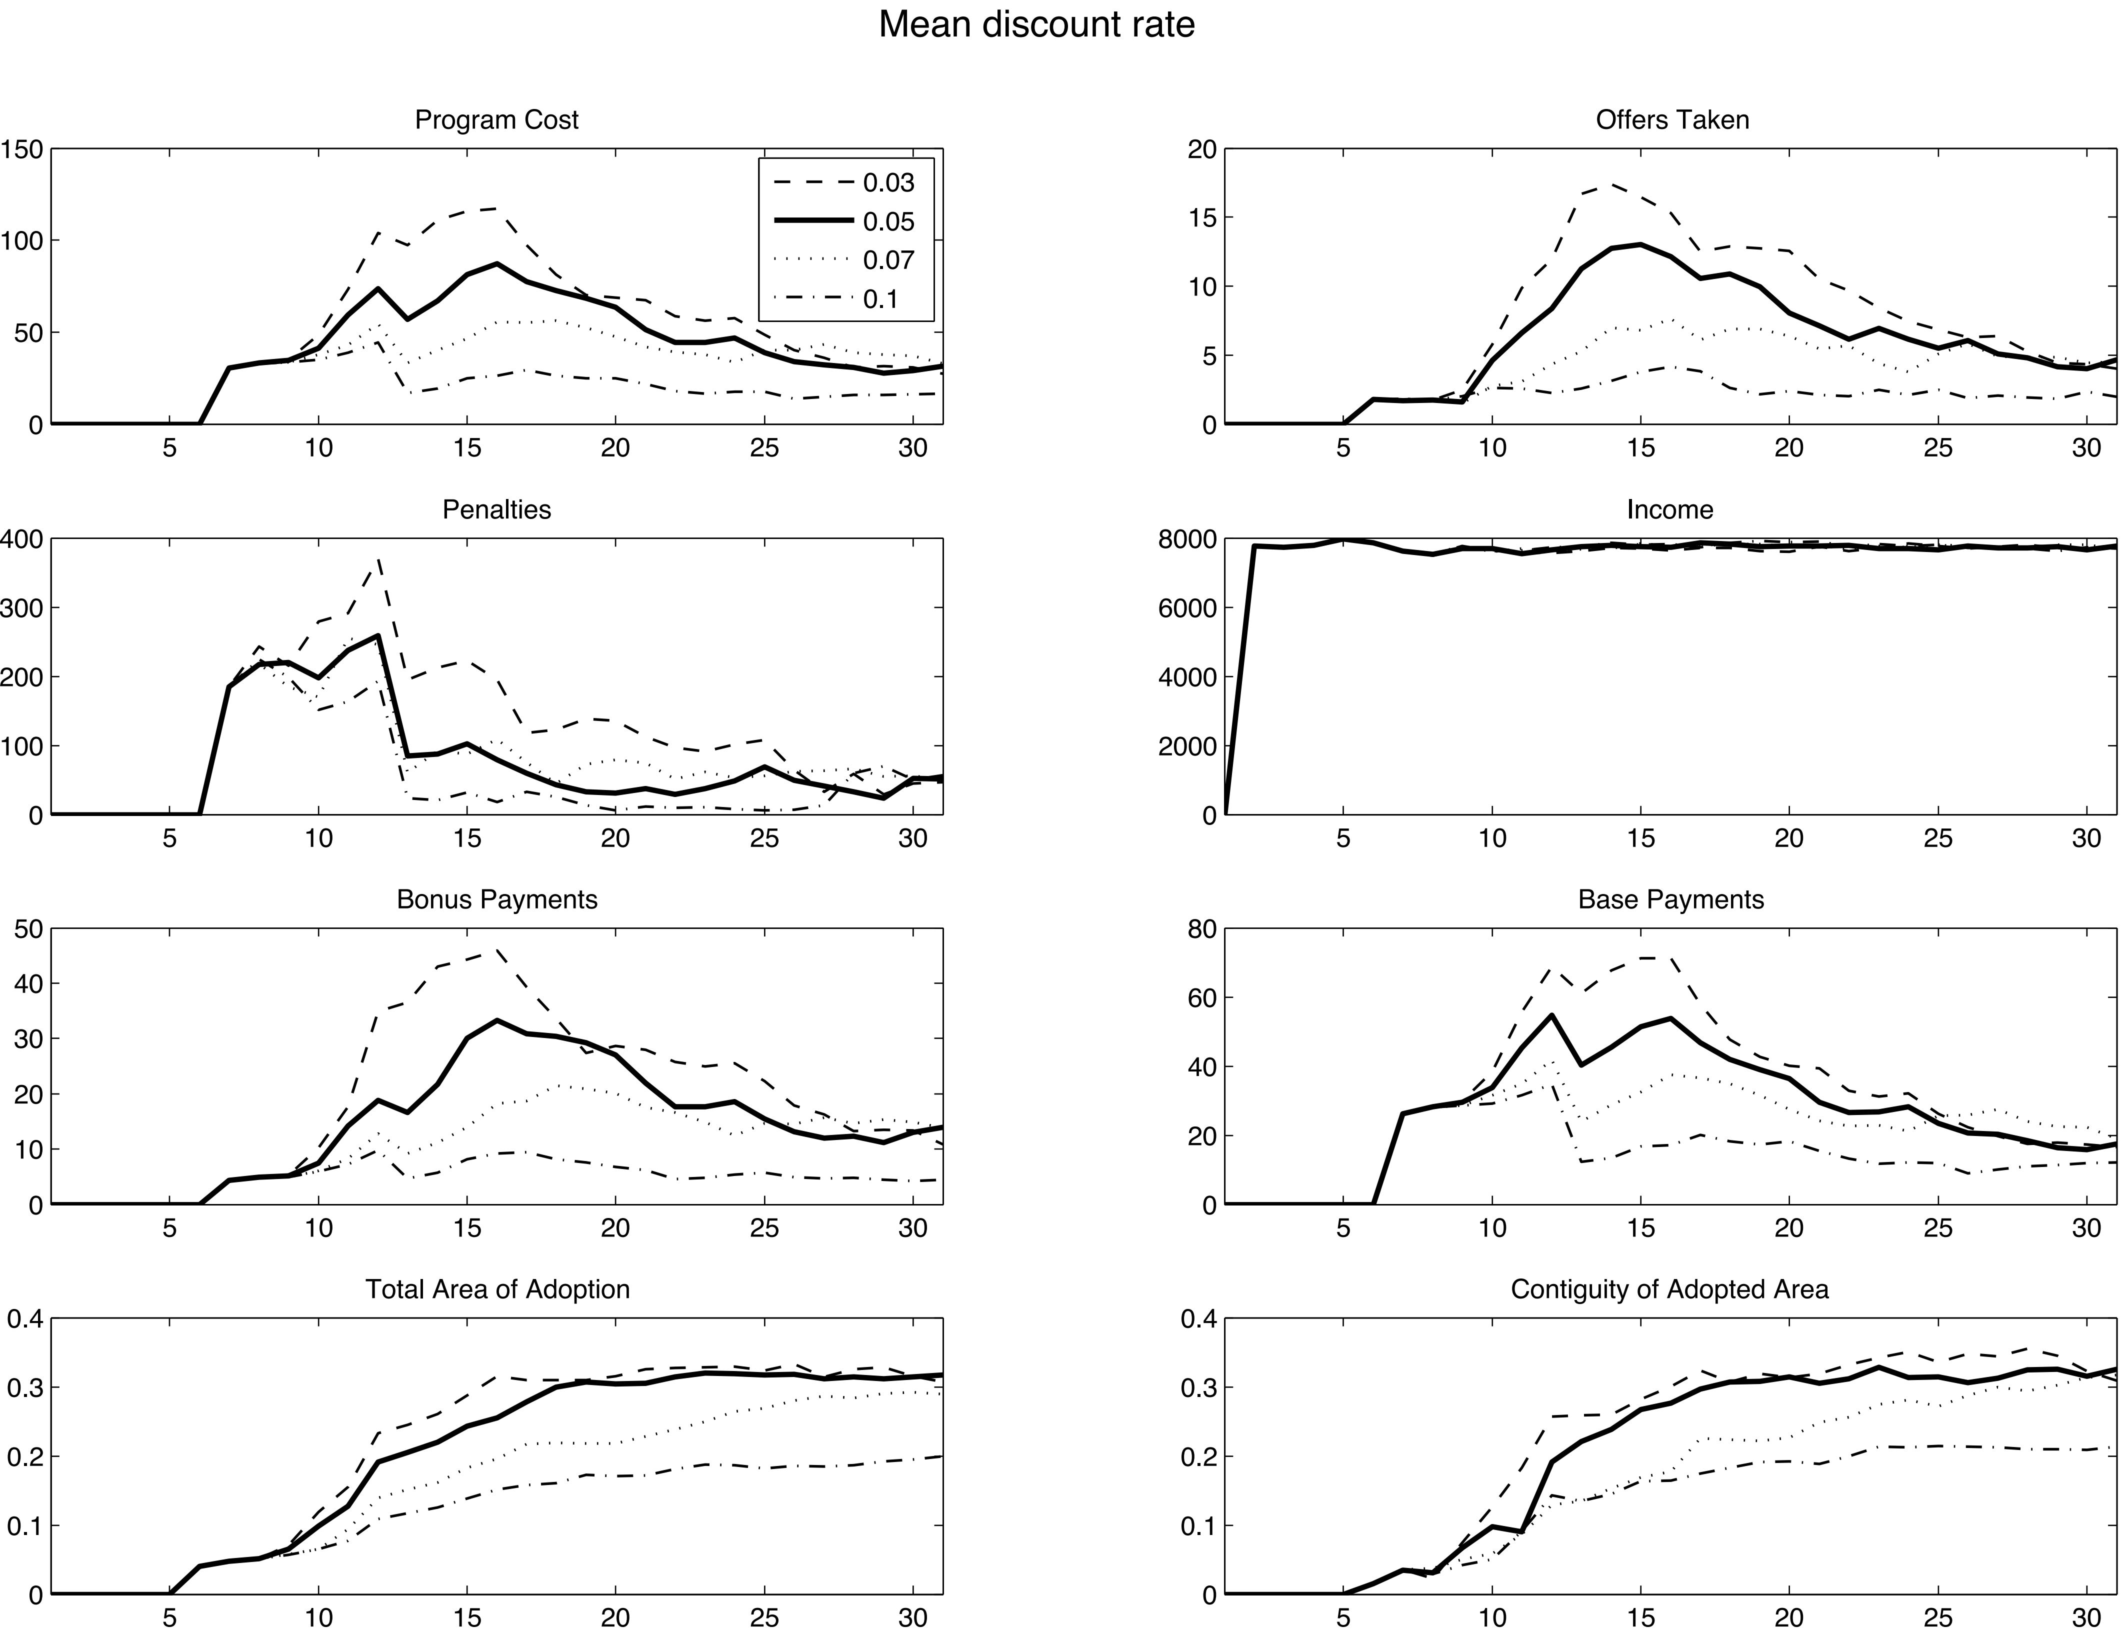

Mean discount rate
Program Cost
0.03
0.05
0.07
0.1
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

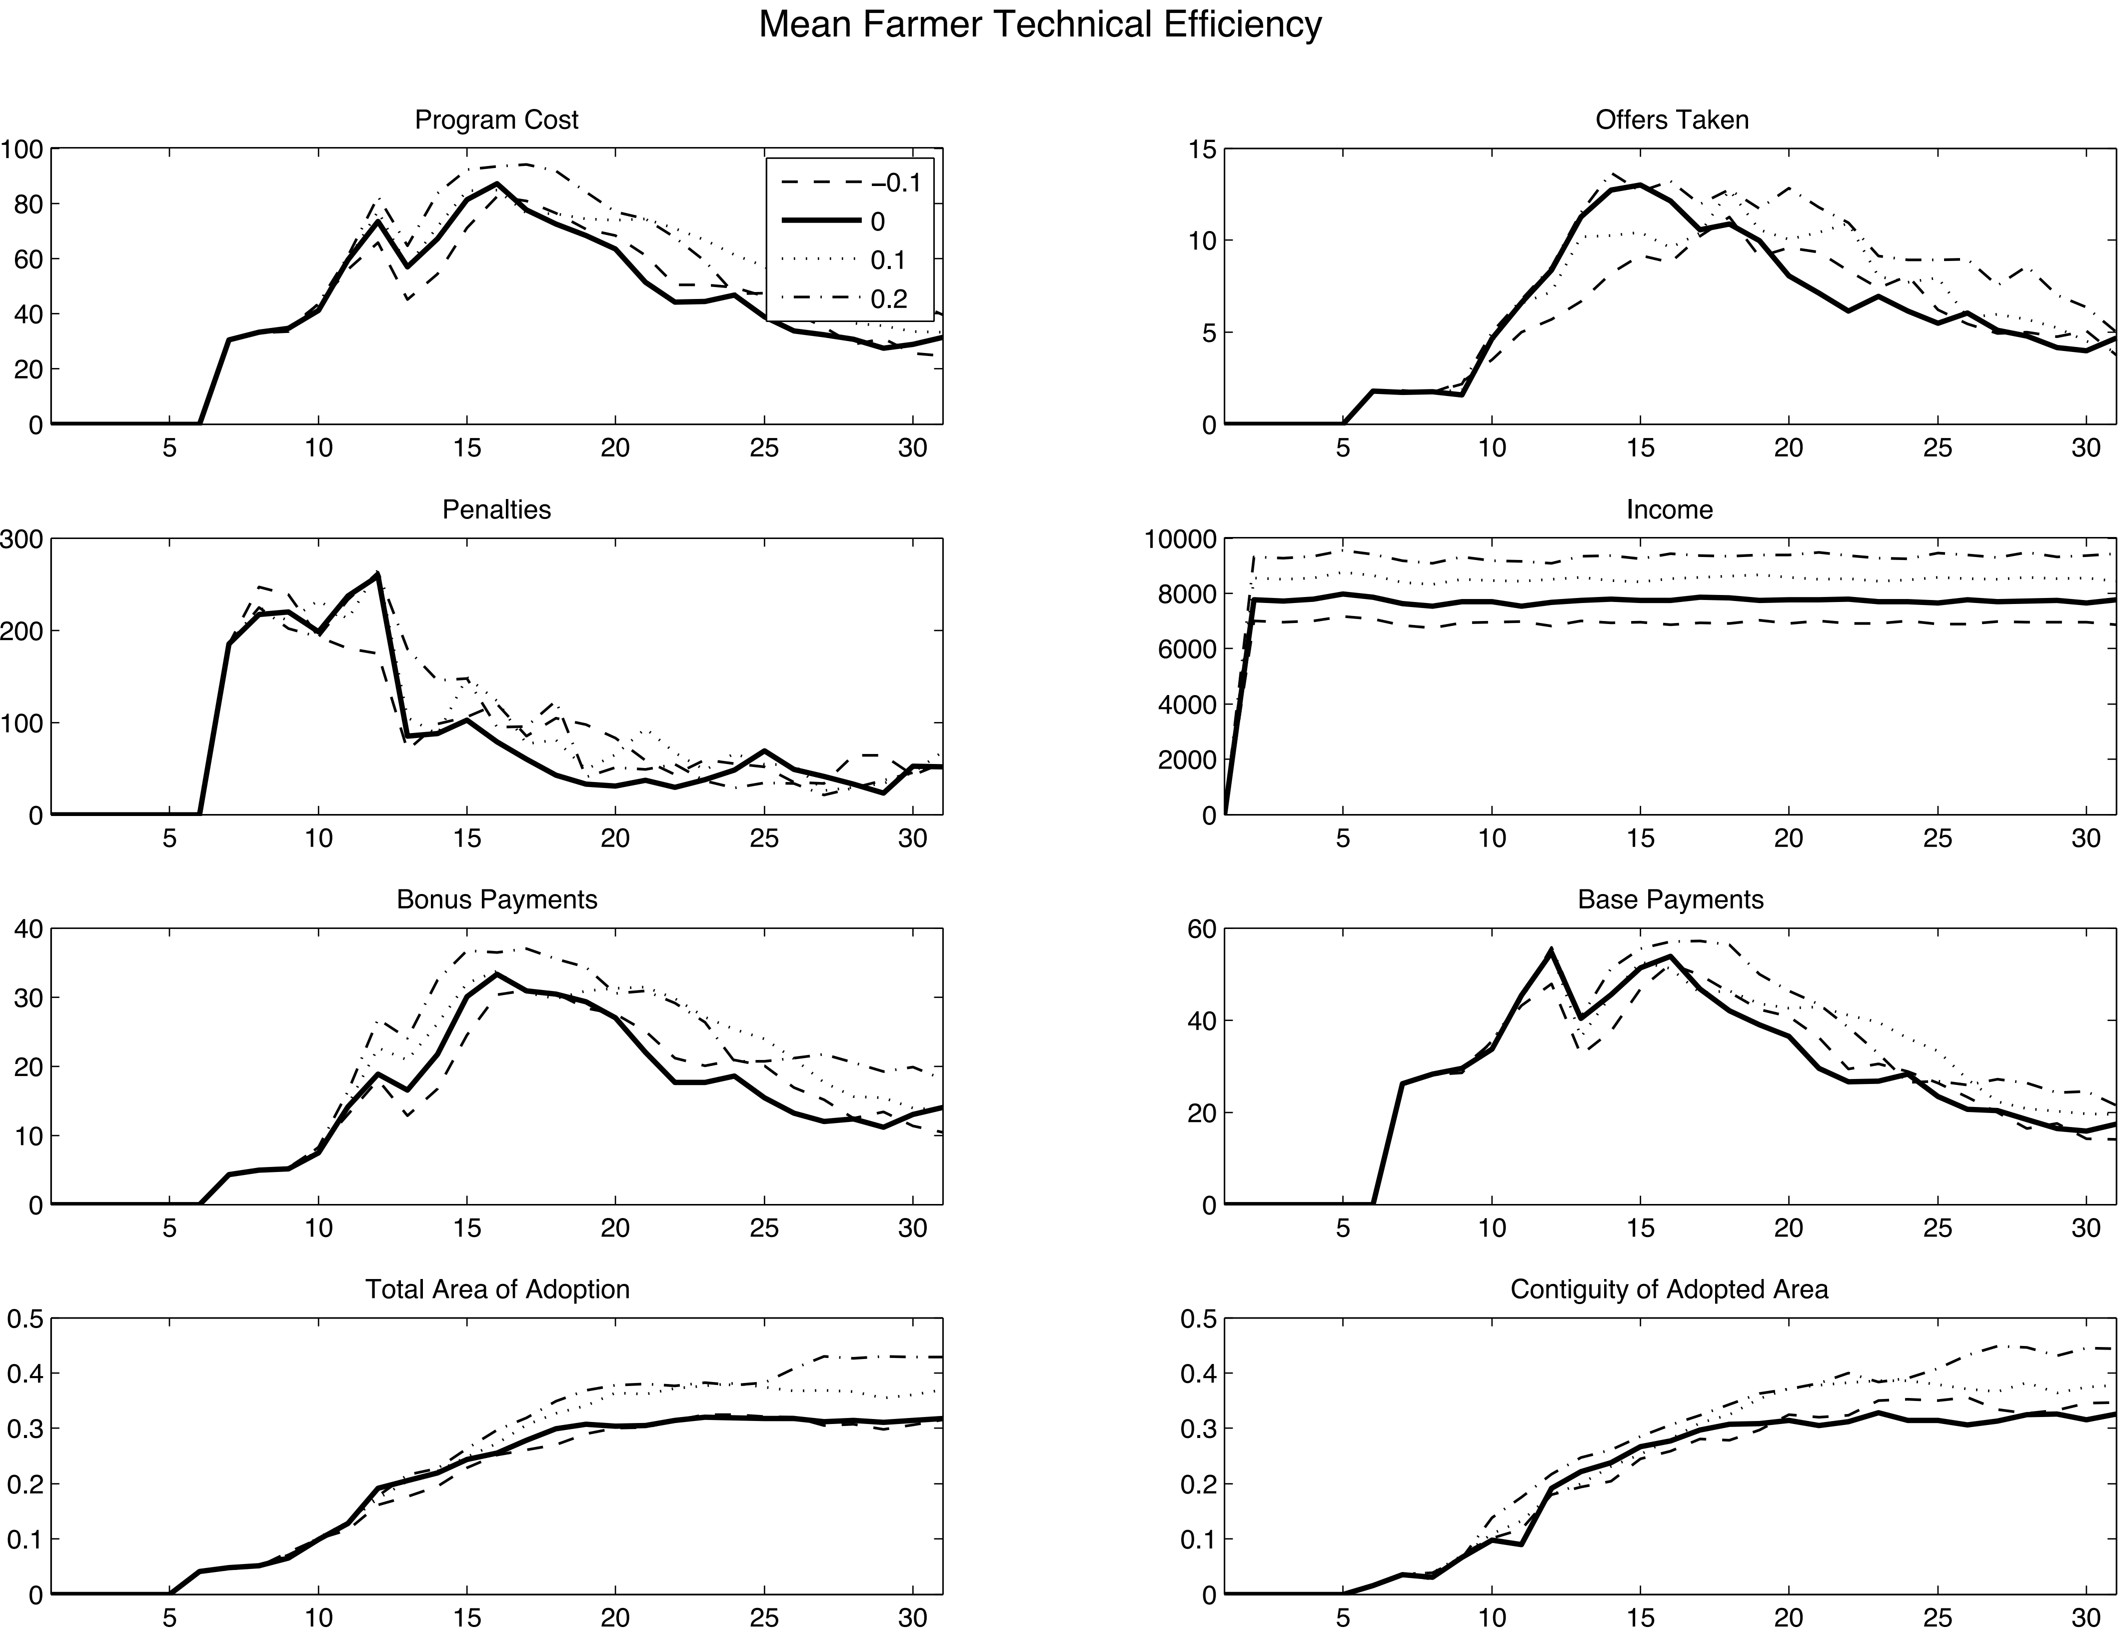

Mean Farmer Technical Efficiency
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
−0.1
0
0.1
0.2

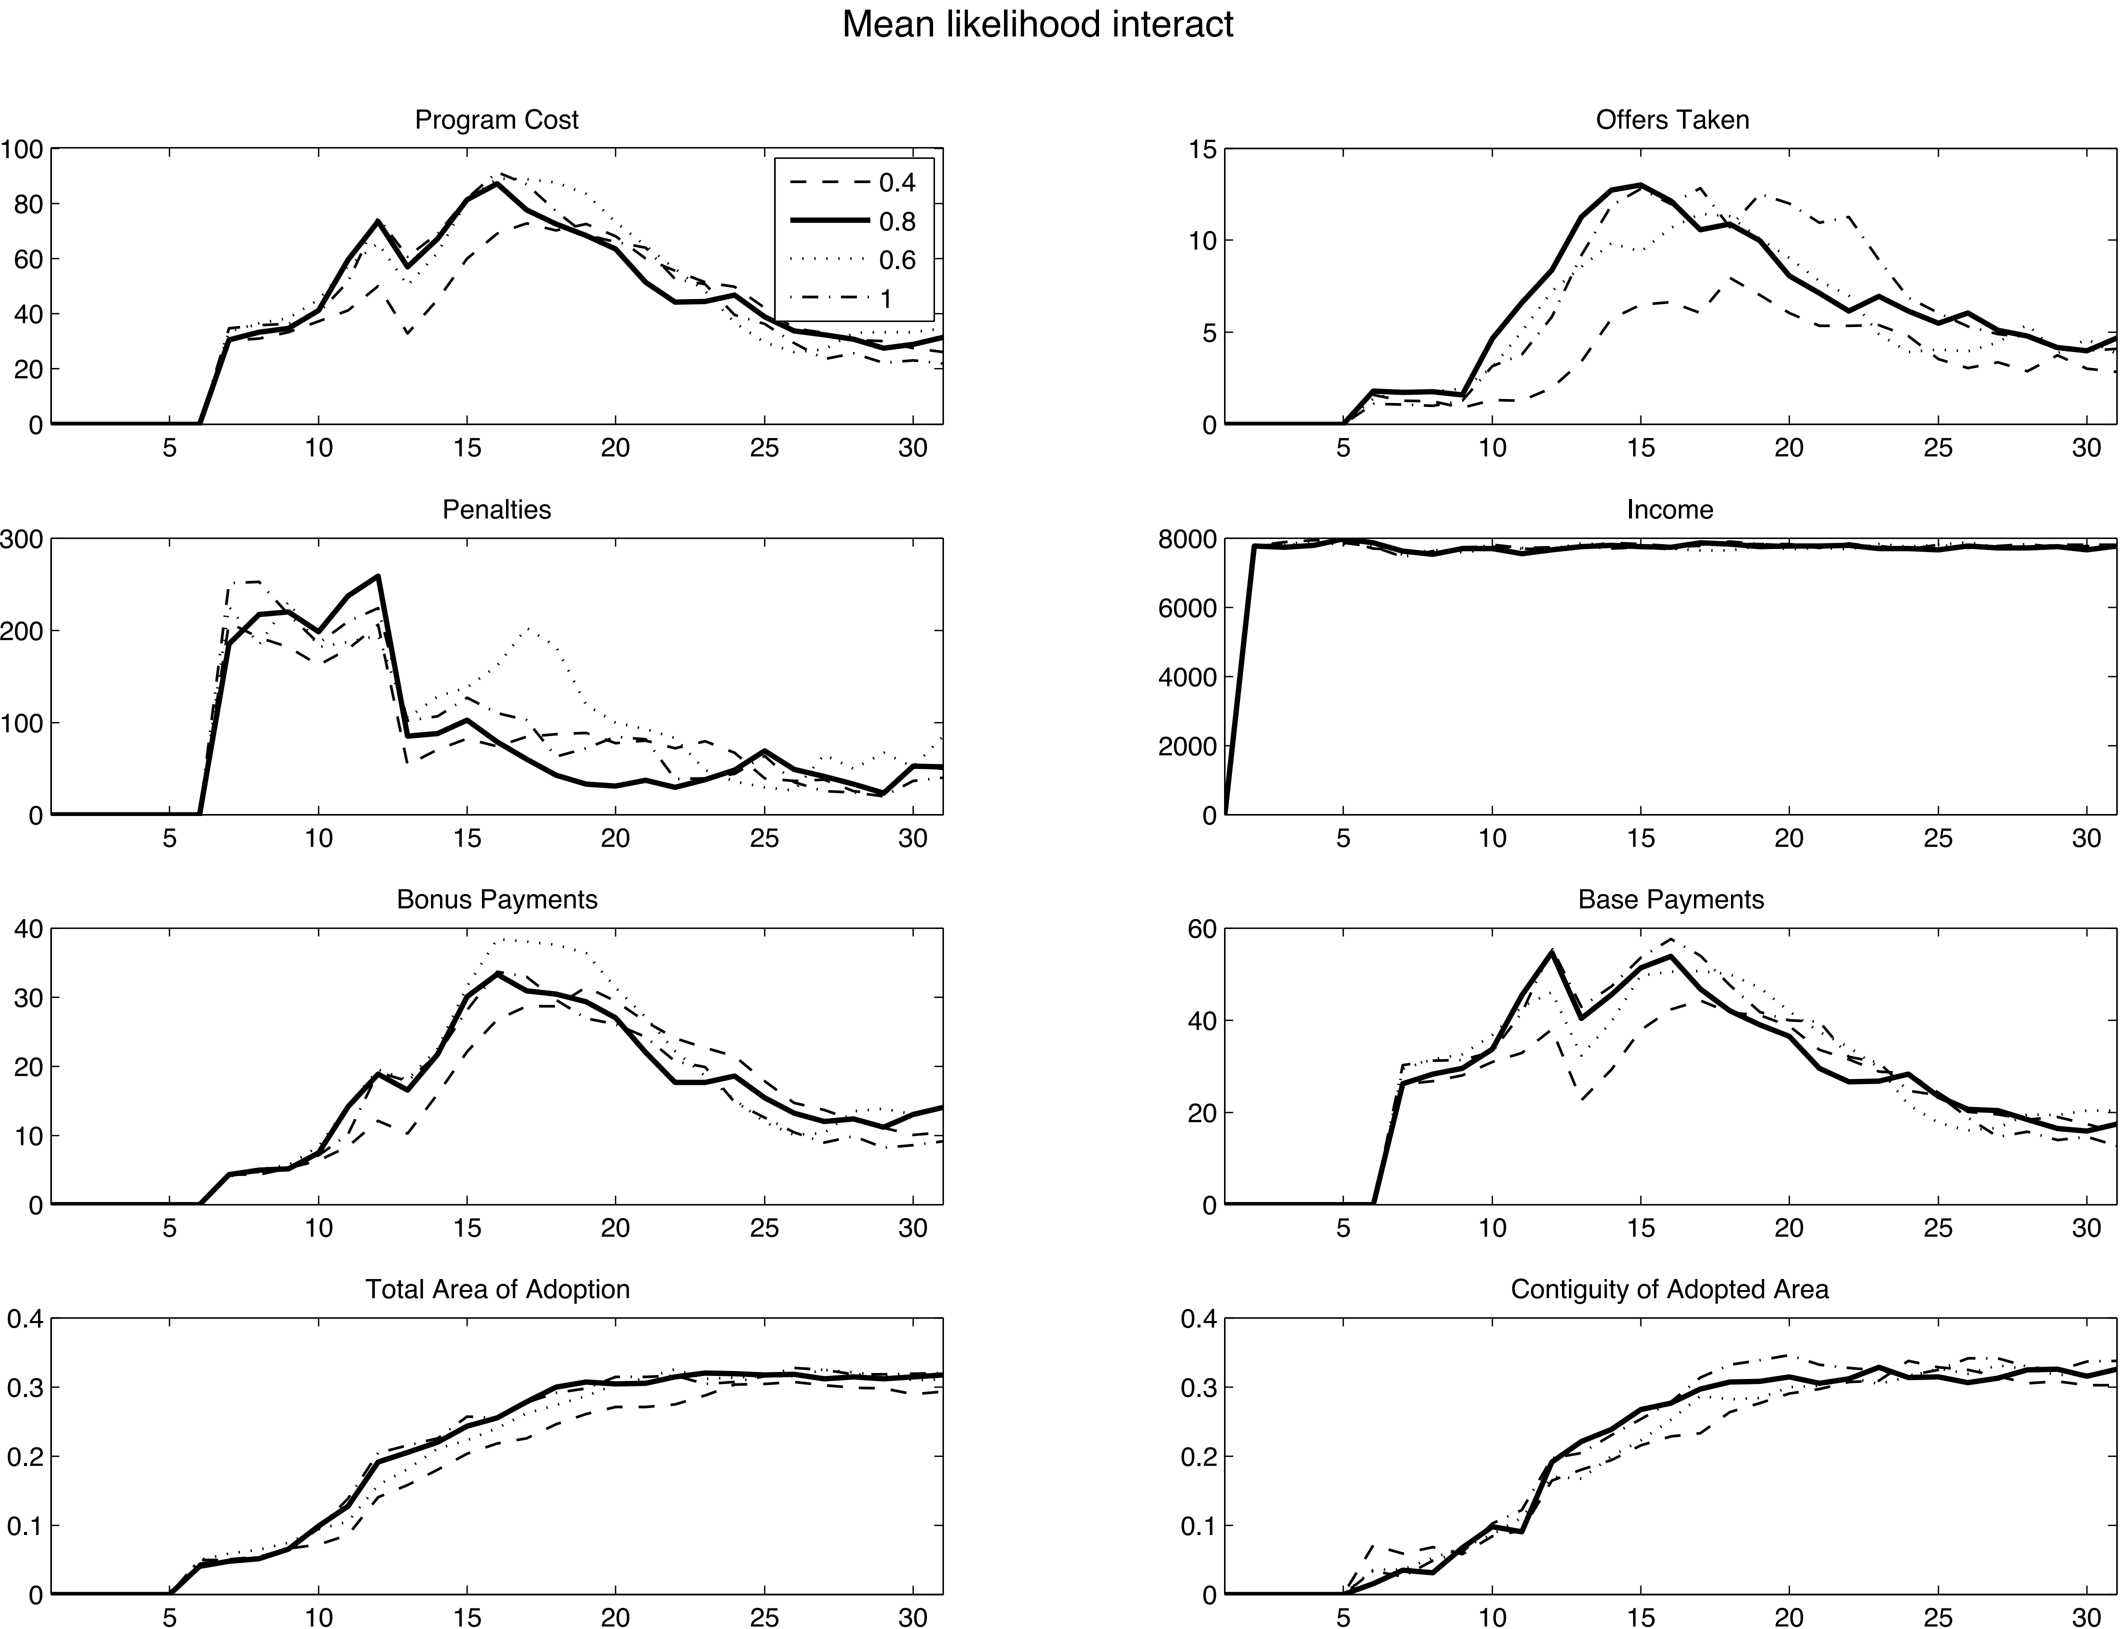

Mean likelihood interact
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
0.4
0.8
0.6
1
0
20
40
60
80
100
5
10
15
25
30
200
300
2000
4000
6000
8000
0.1
0.2
0.3

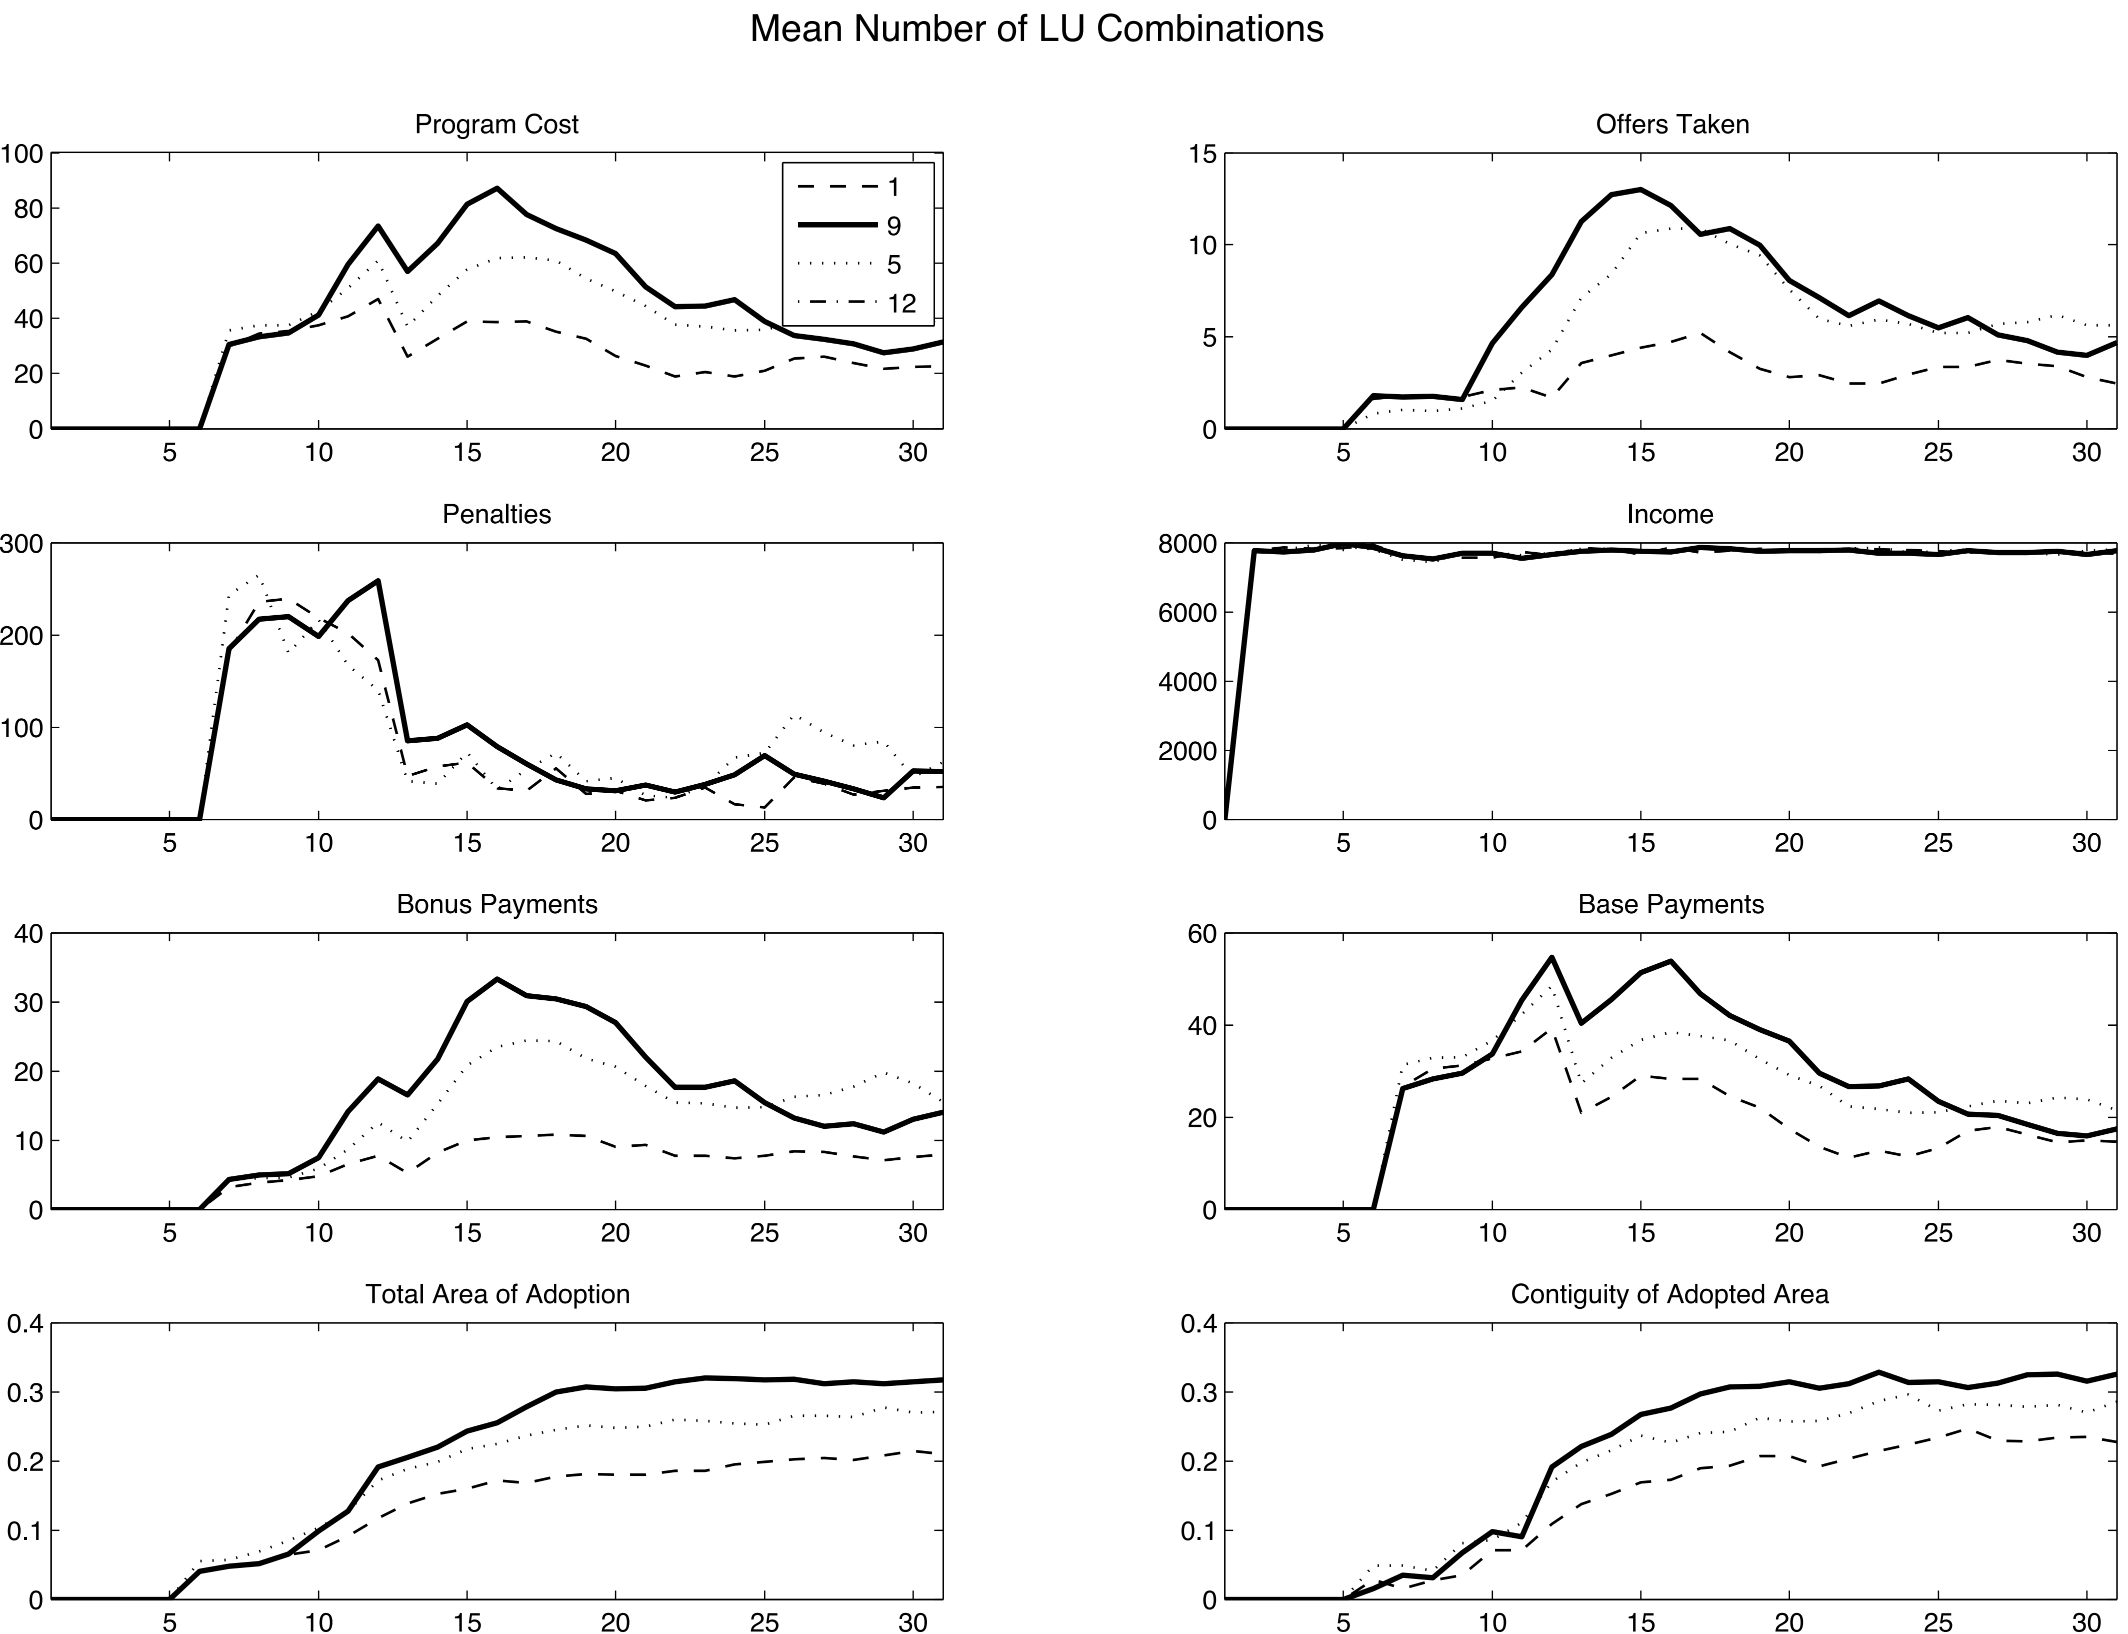

Mean Number of LU Combinations
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
1
9
5
12

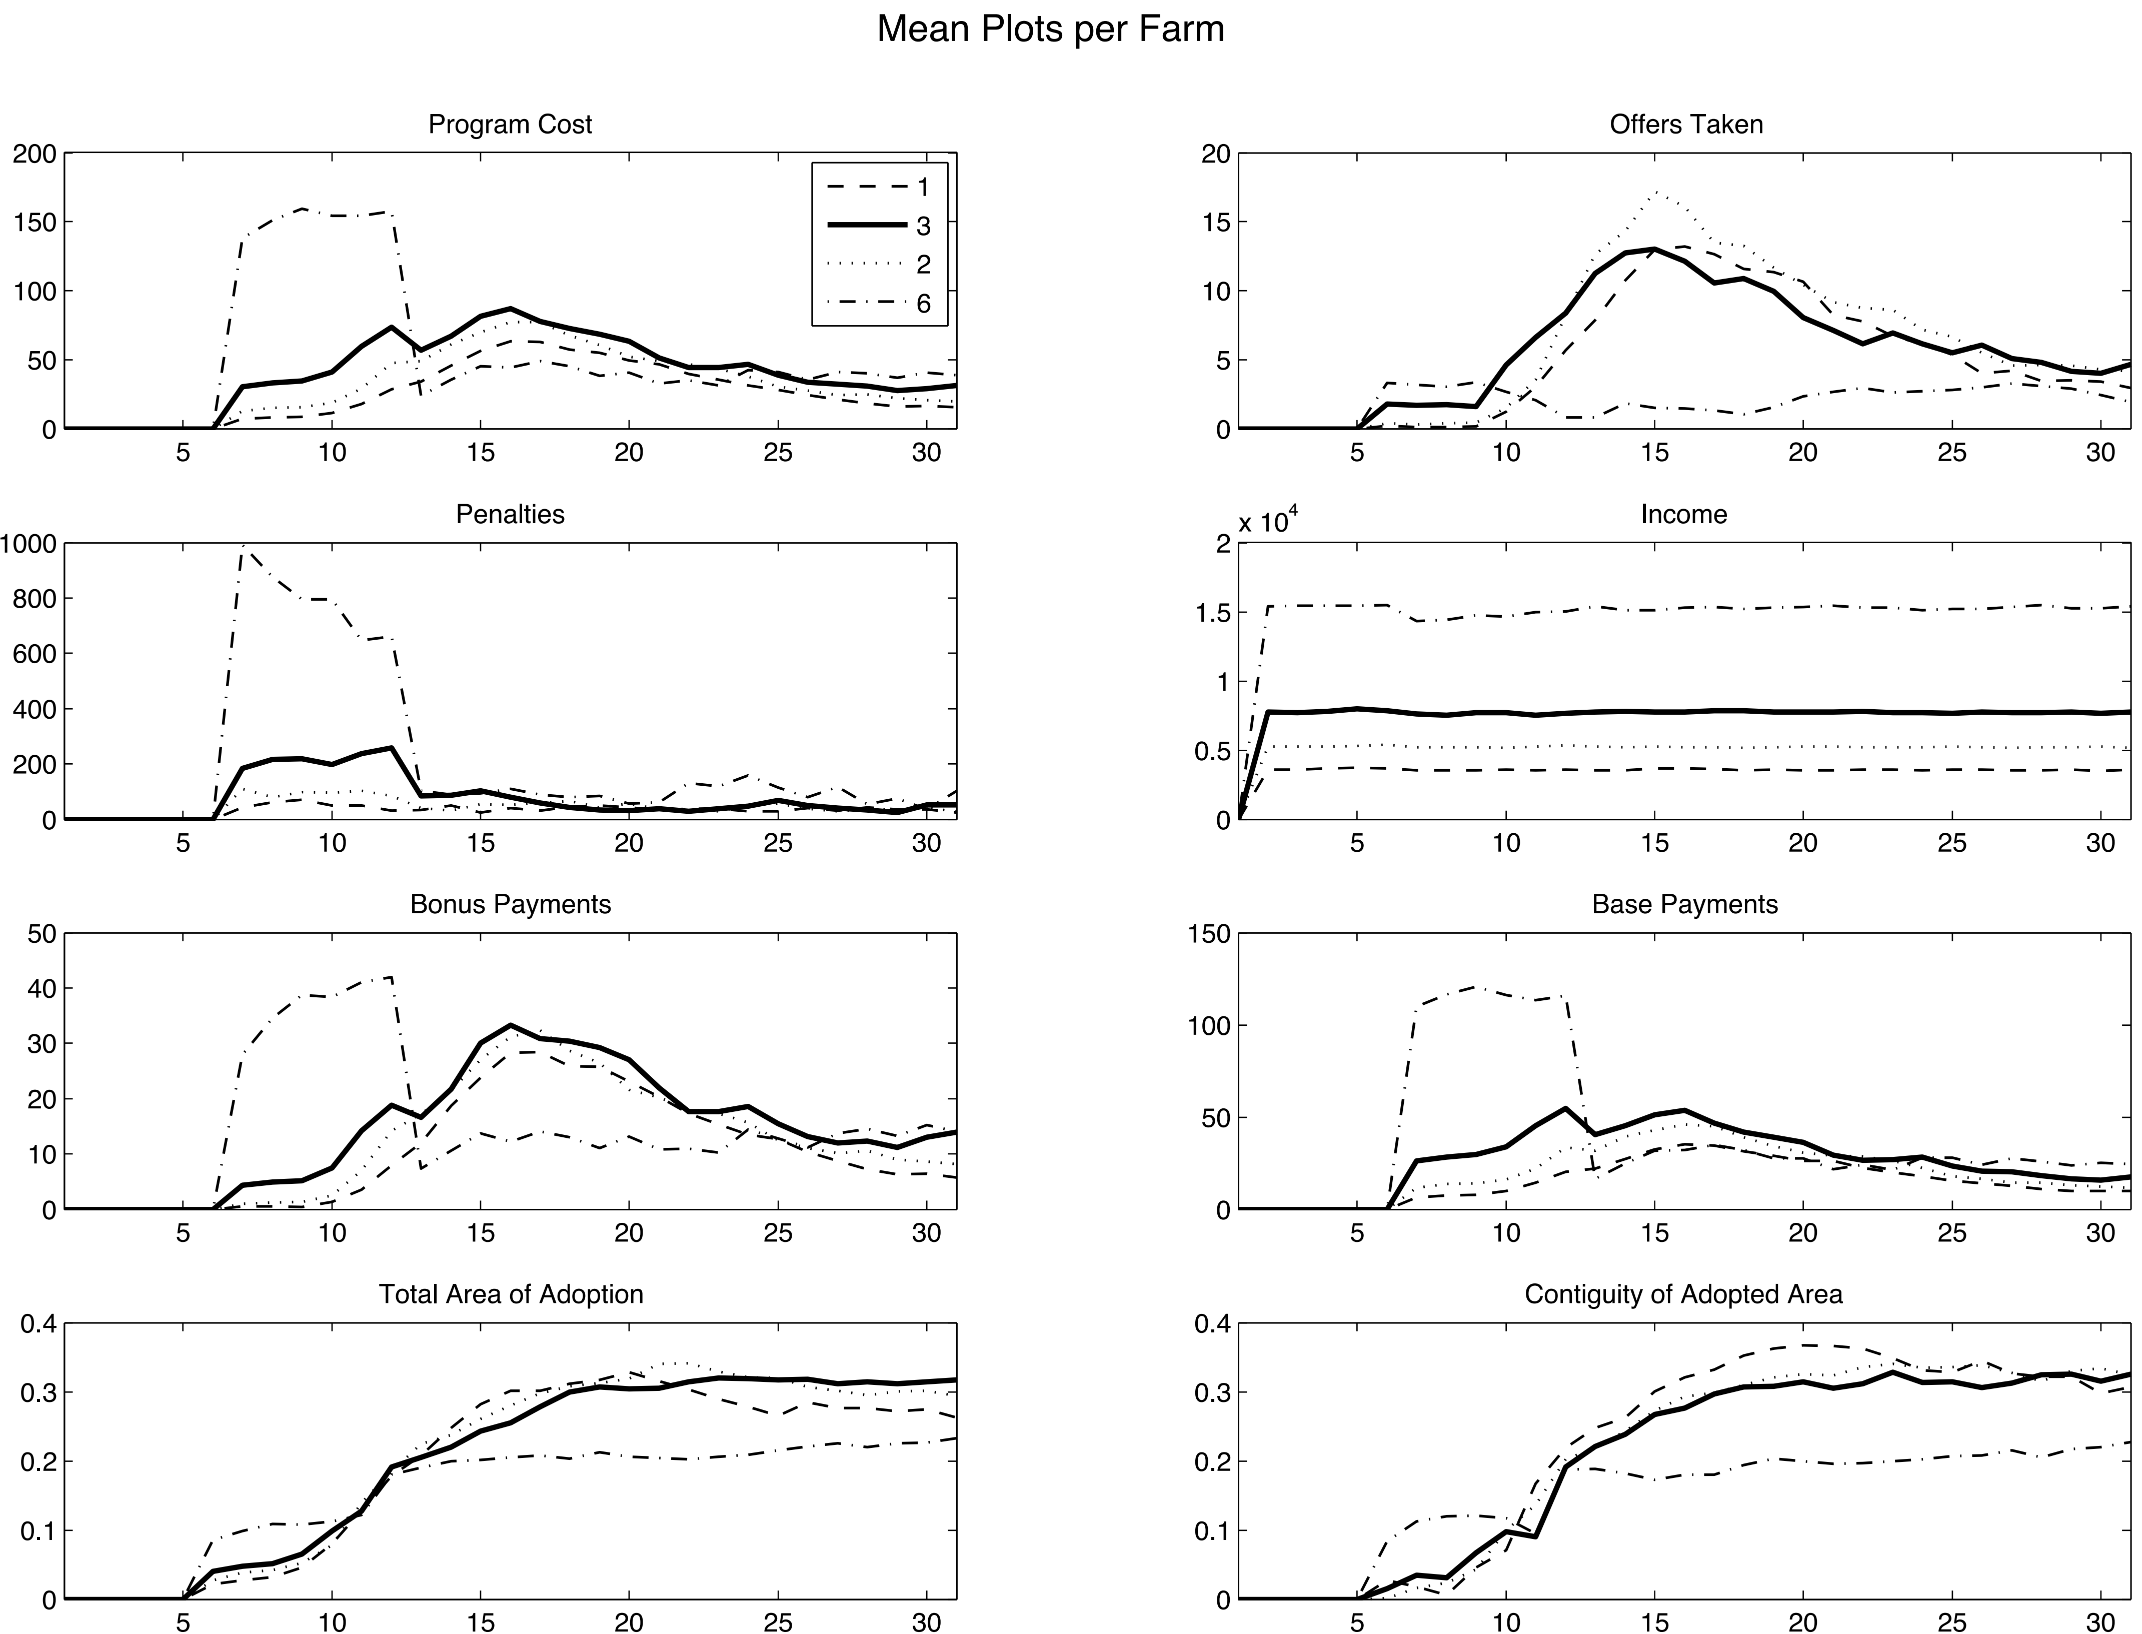

Mean Plots per Farm
Program Cost
Offers Taken
Penalties
Income
x 10^4
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
1
3
2
6

Mean property size

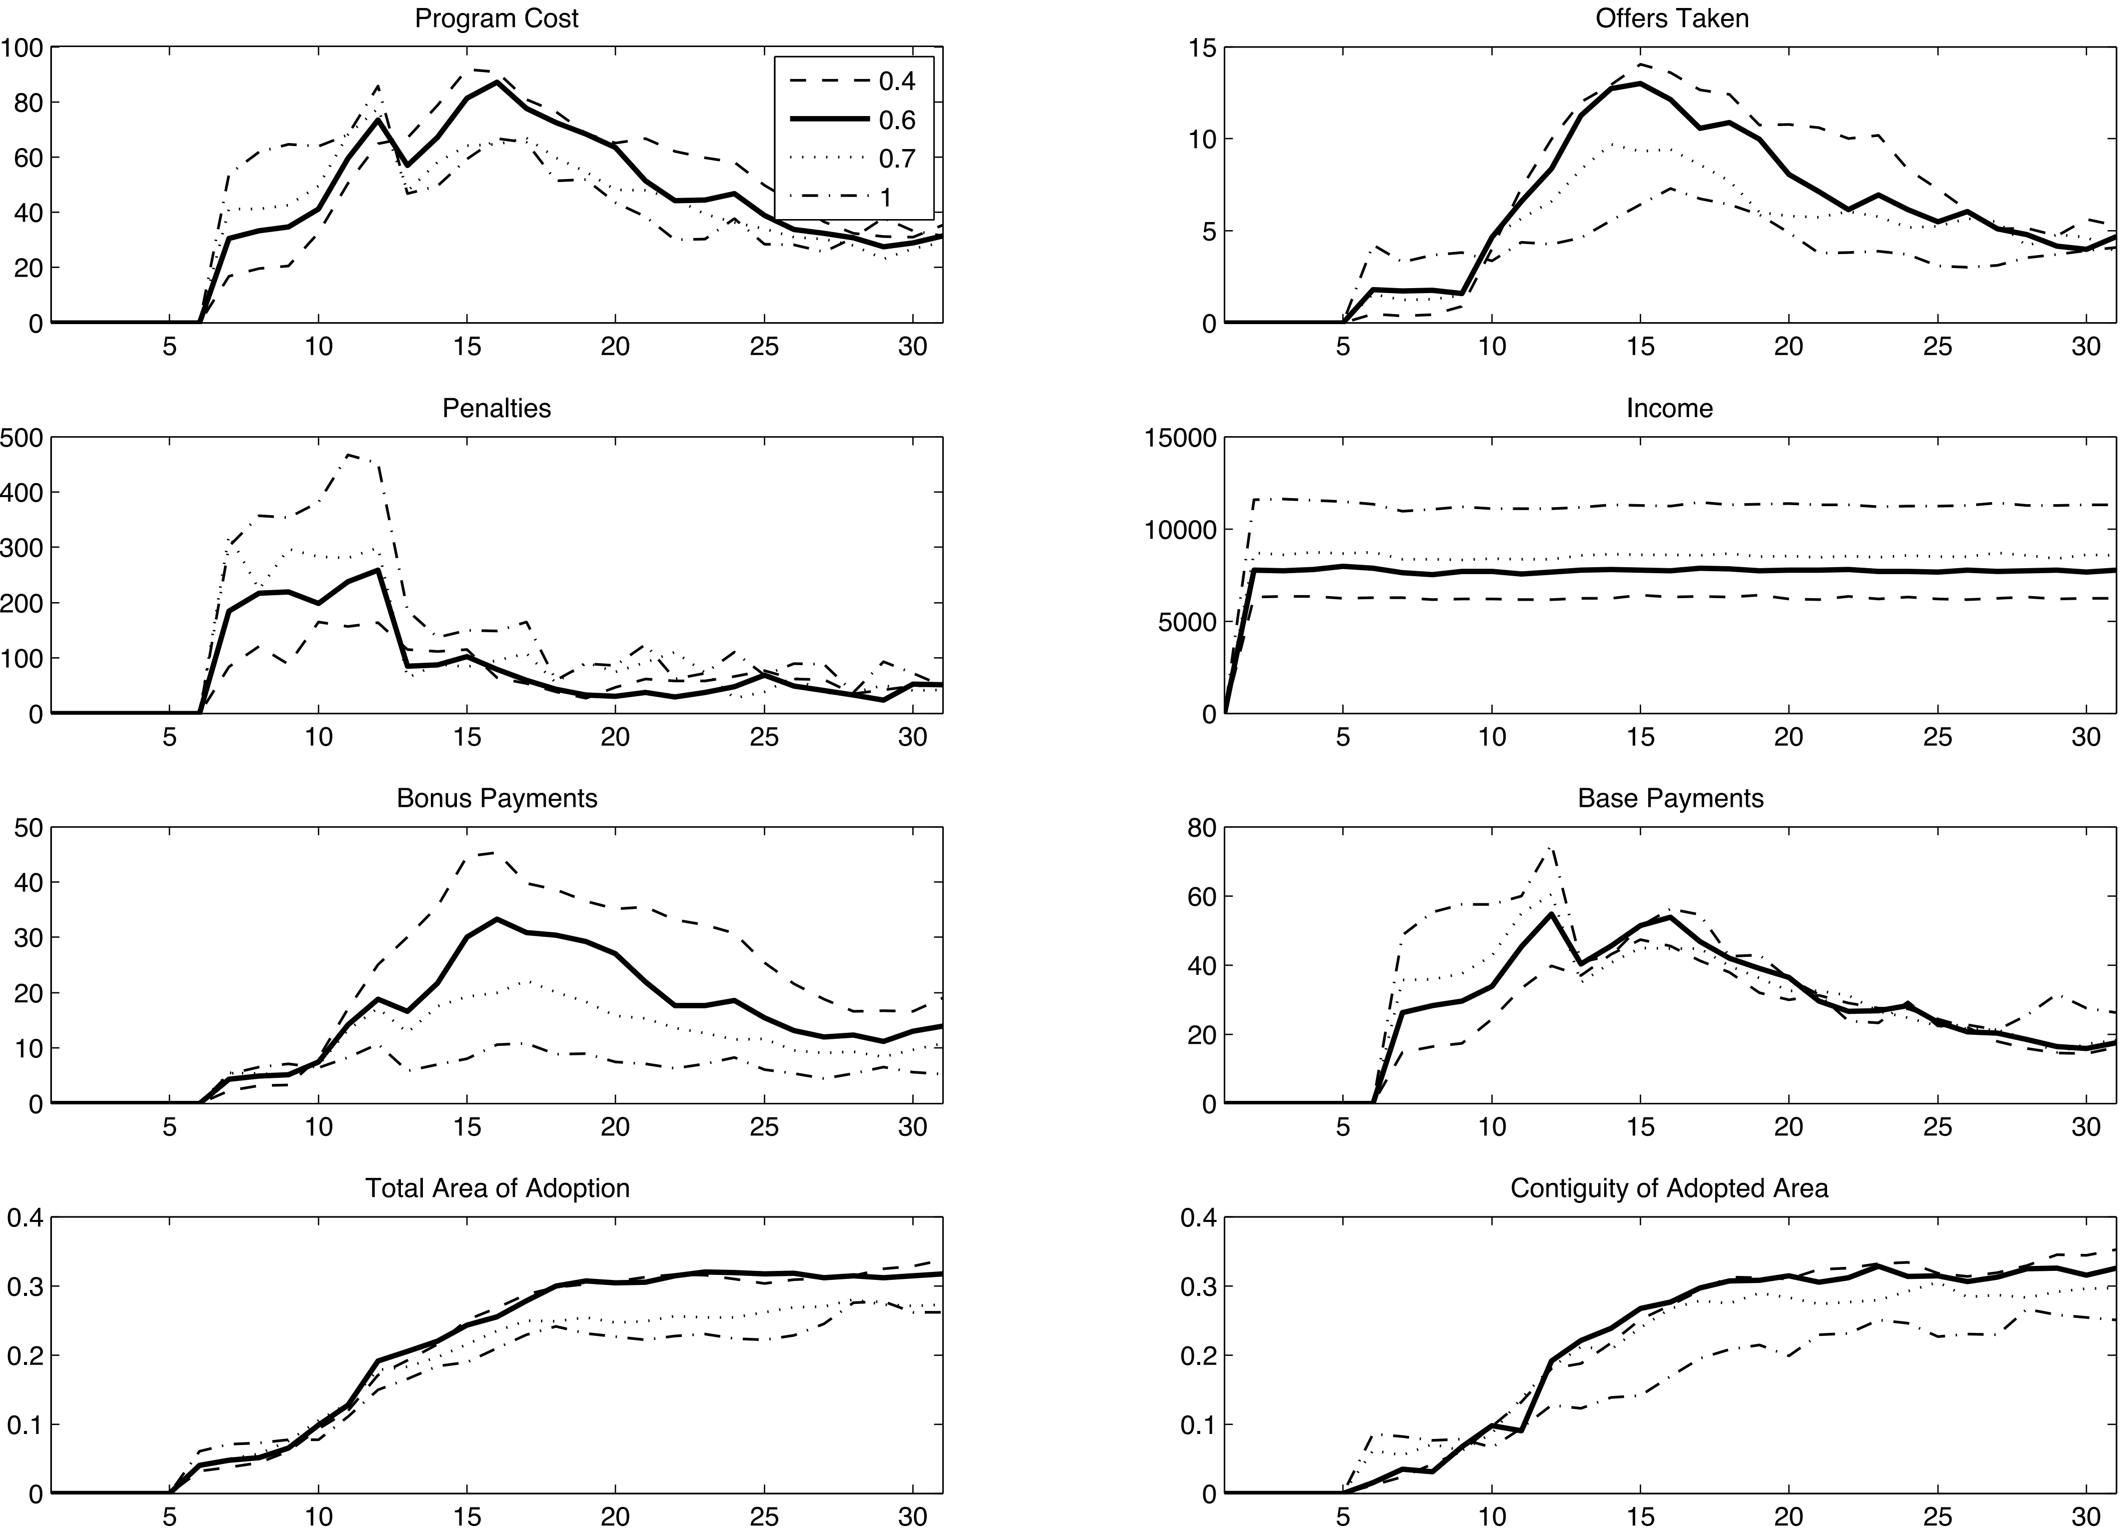

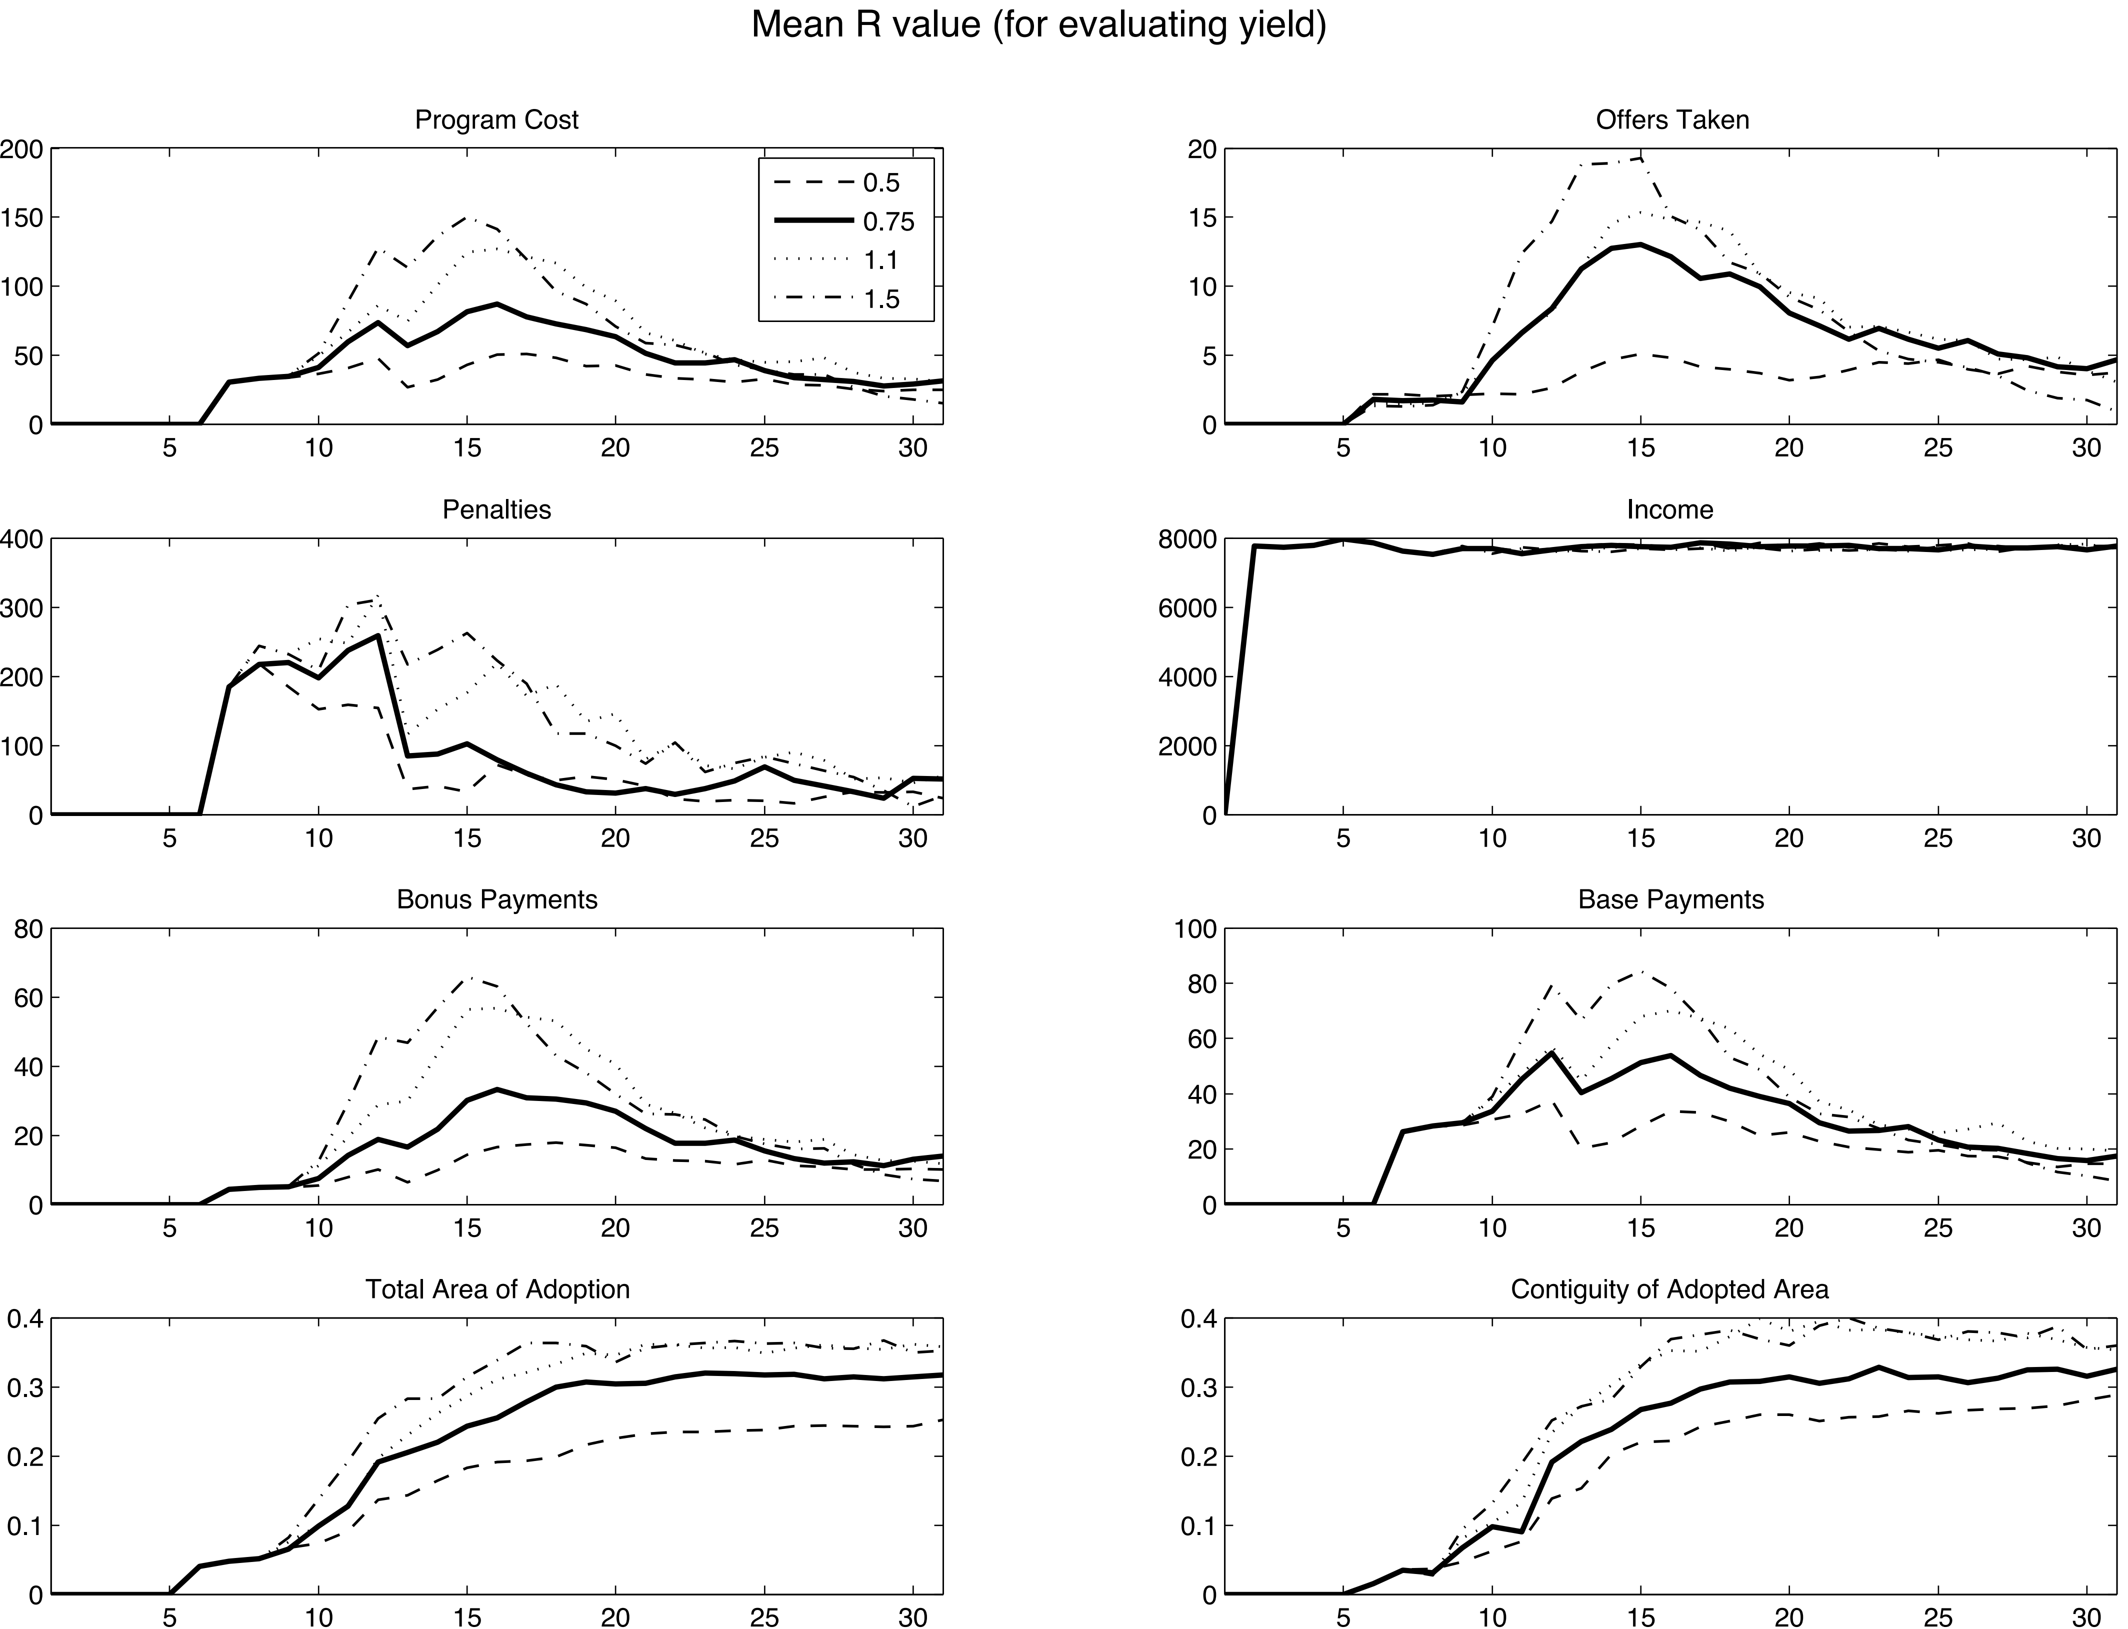

Mean R value (for evaluating yield)
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
0.5
0.75
1.1
1.5

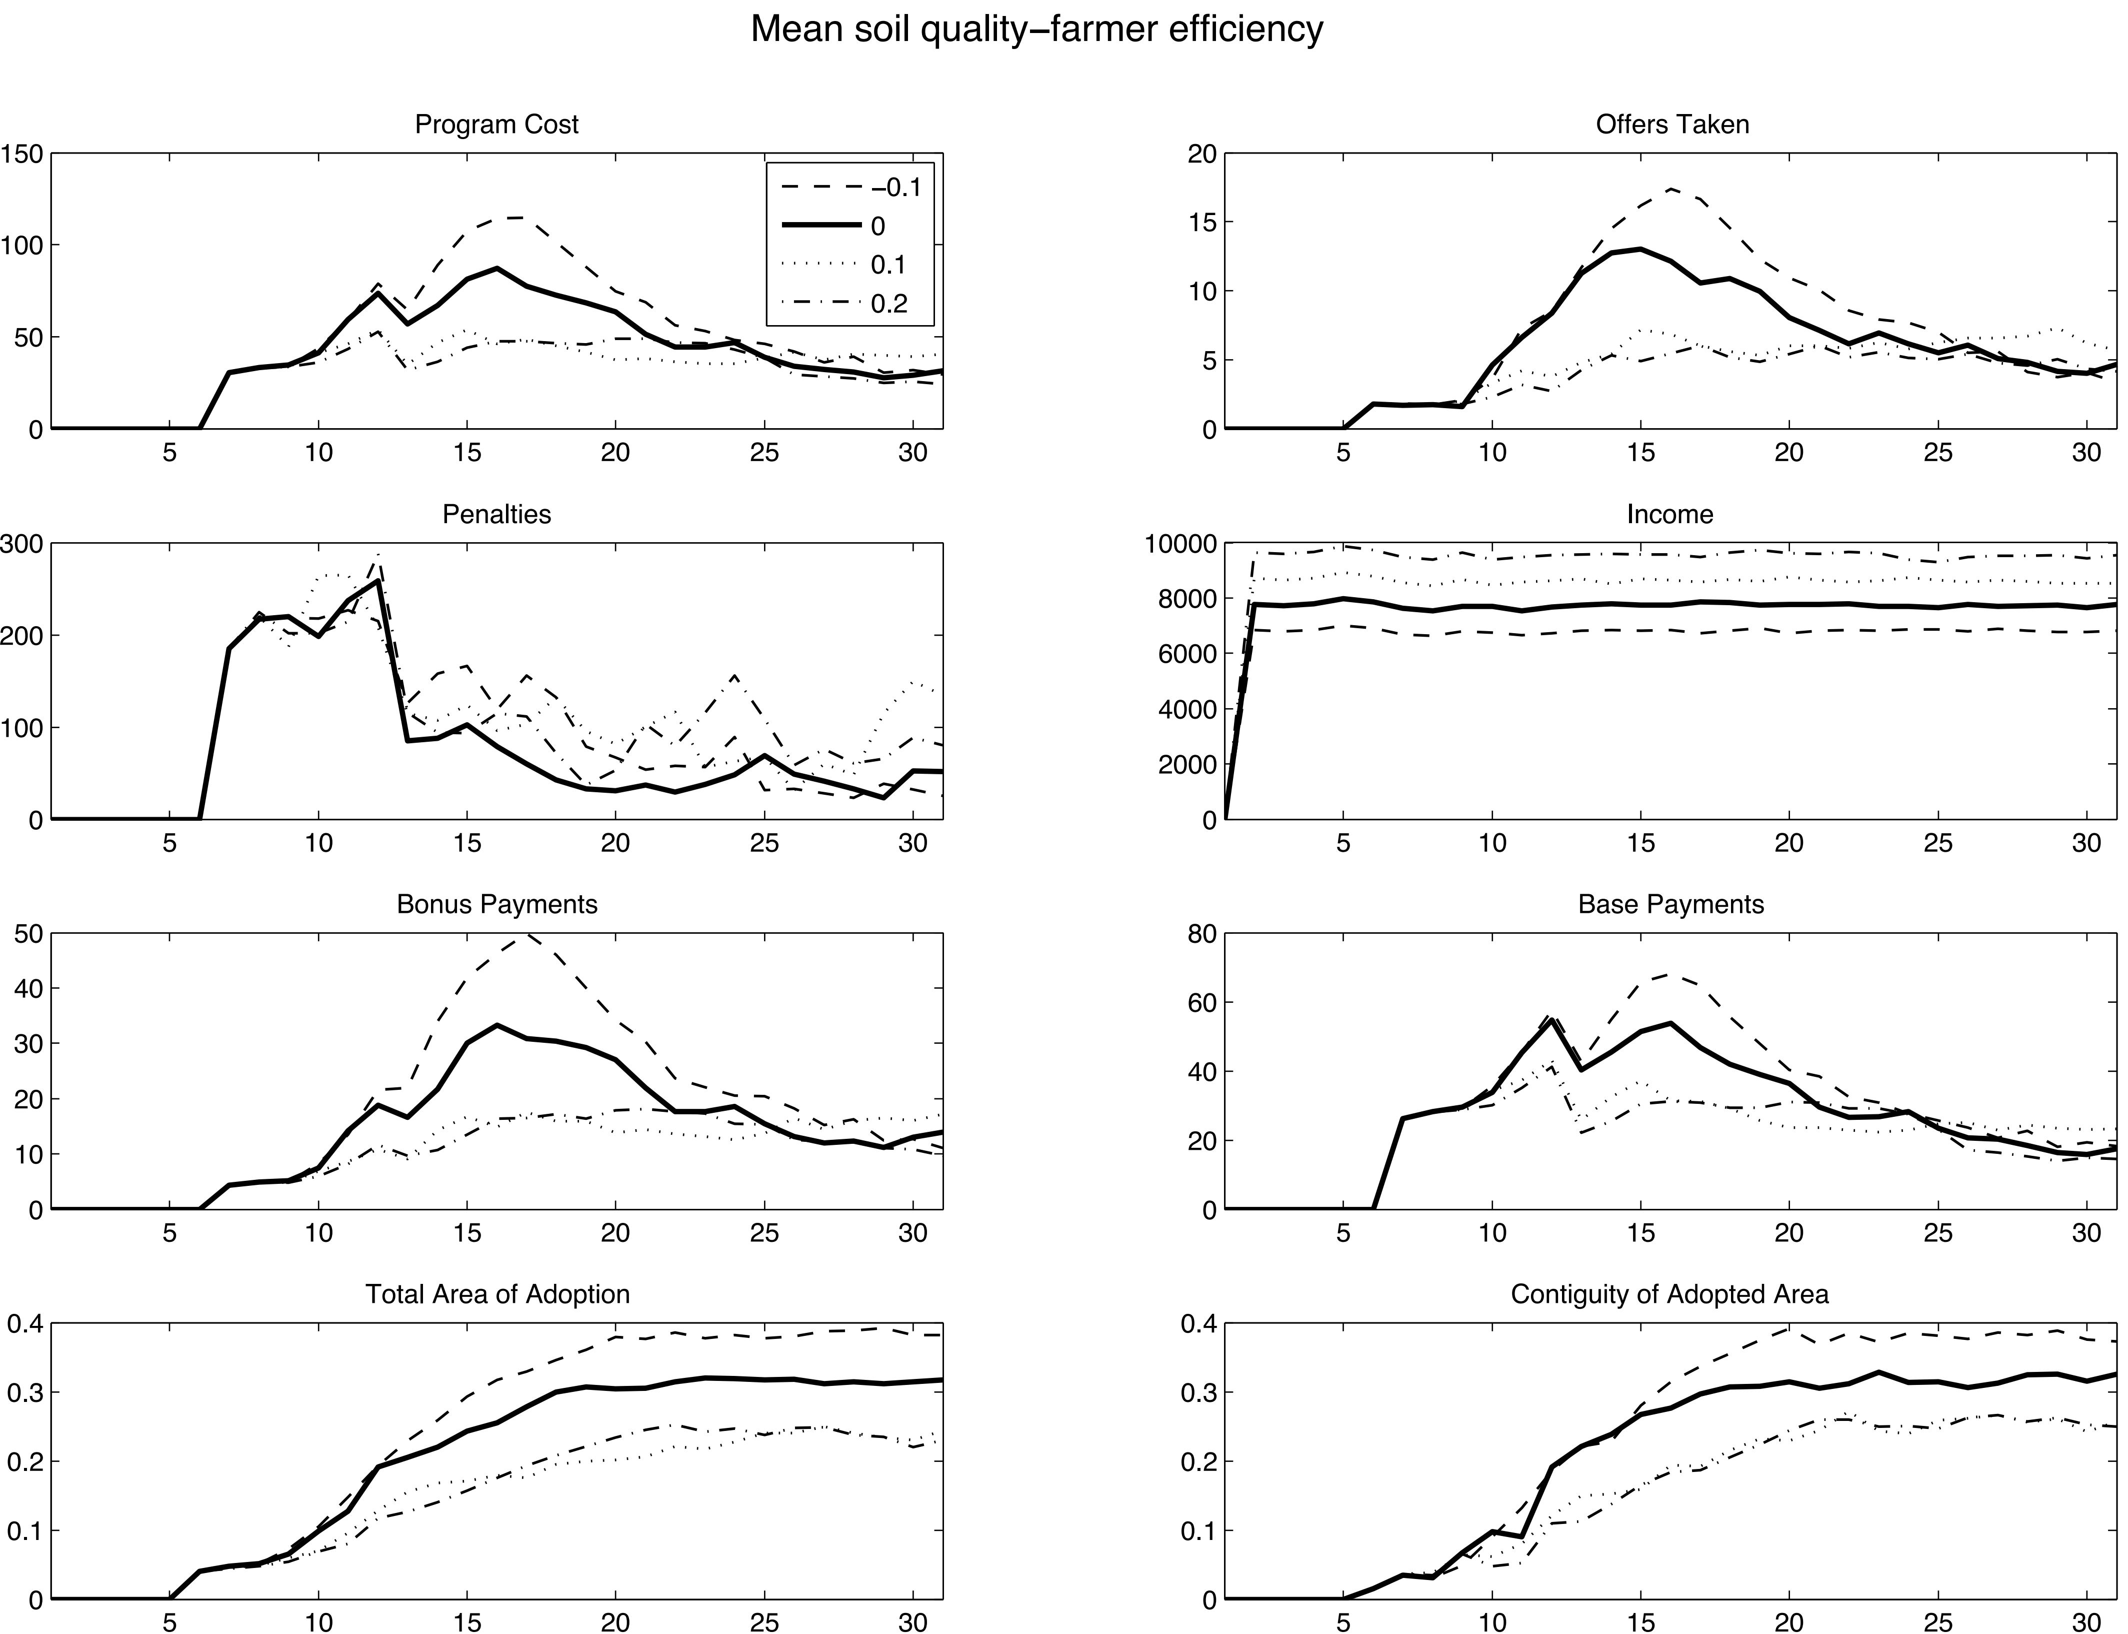

Mean soil quality–farmer efficiency
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
–0.1
0
0.1
0.2

Neighborhood radius

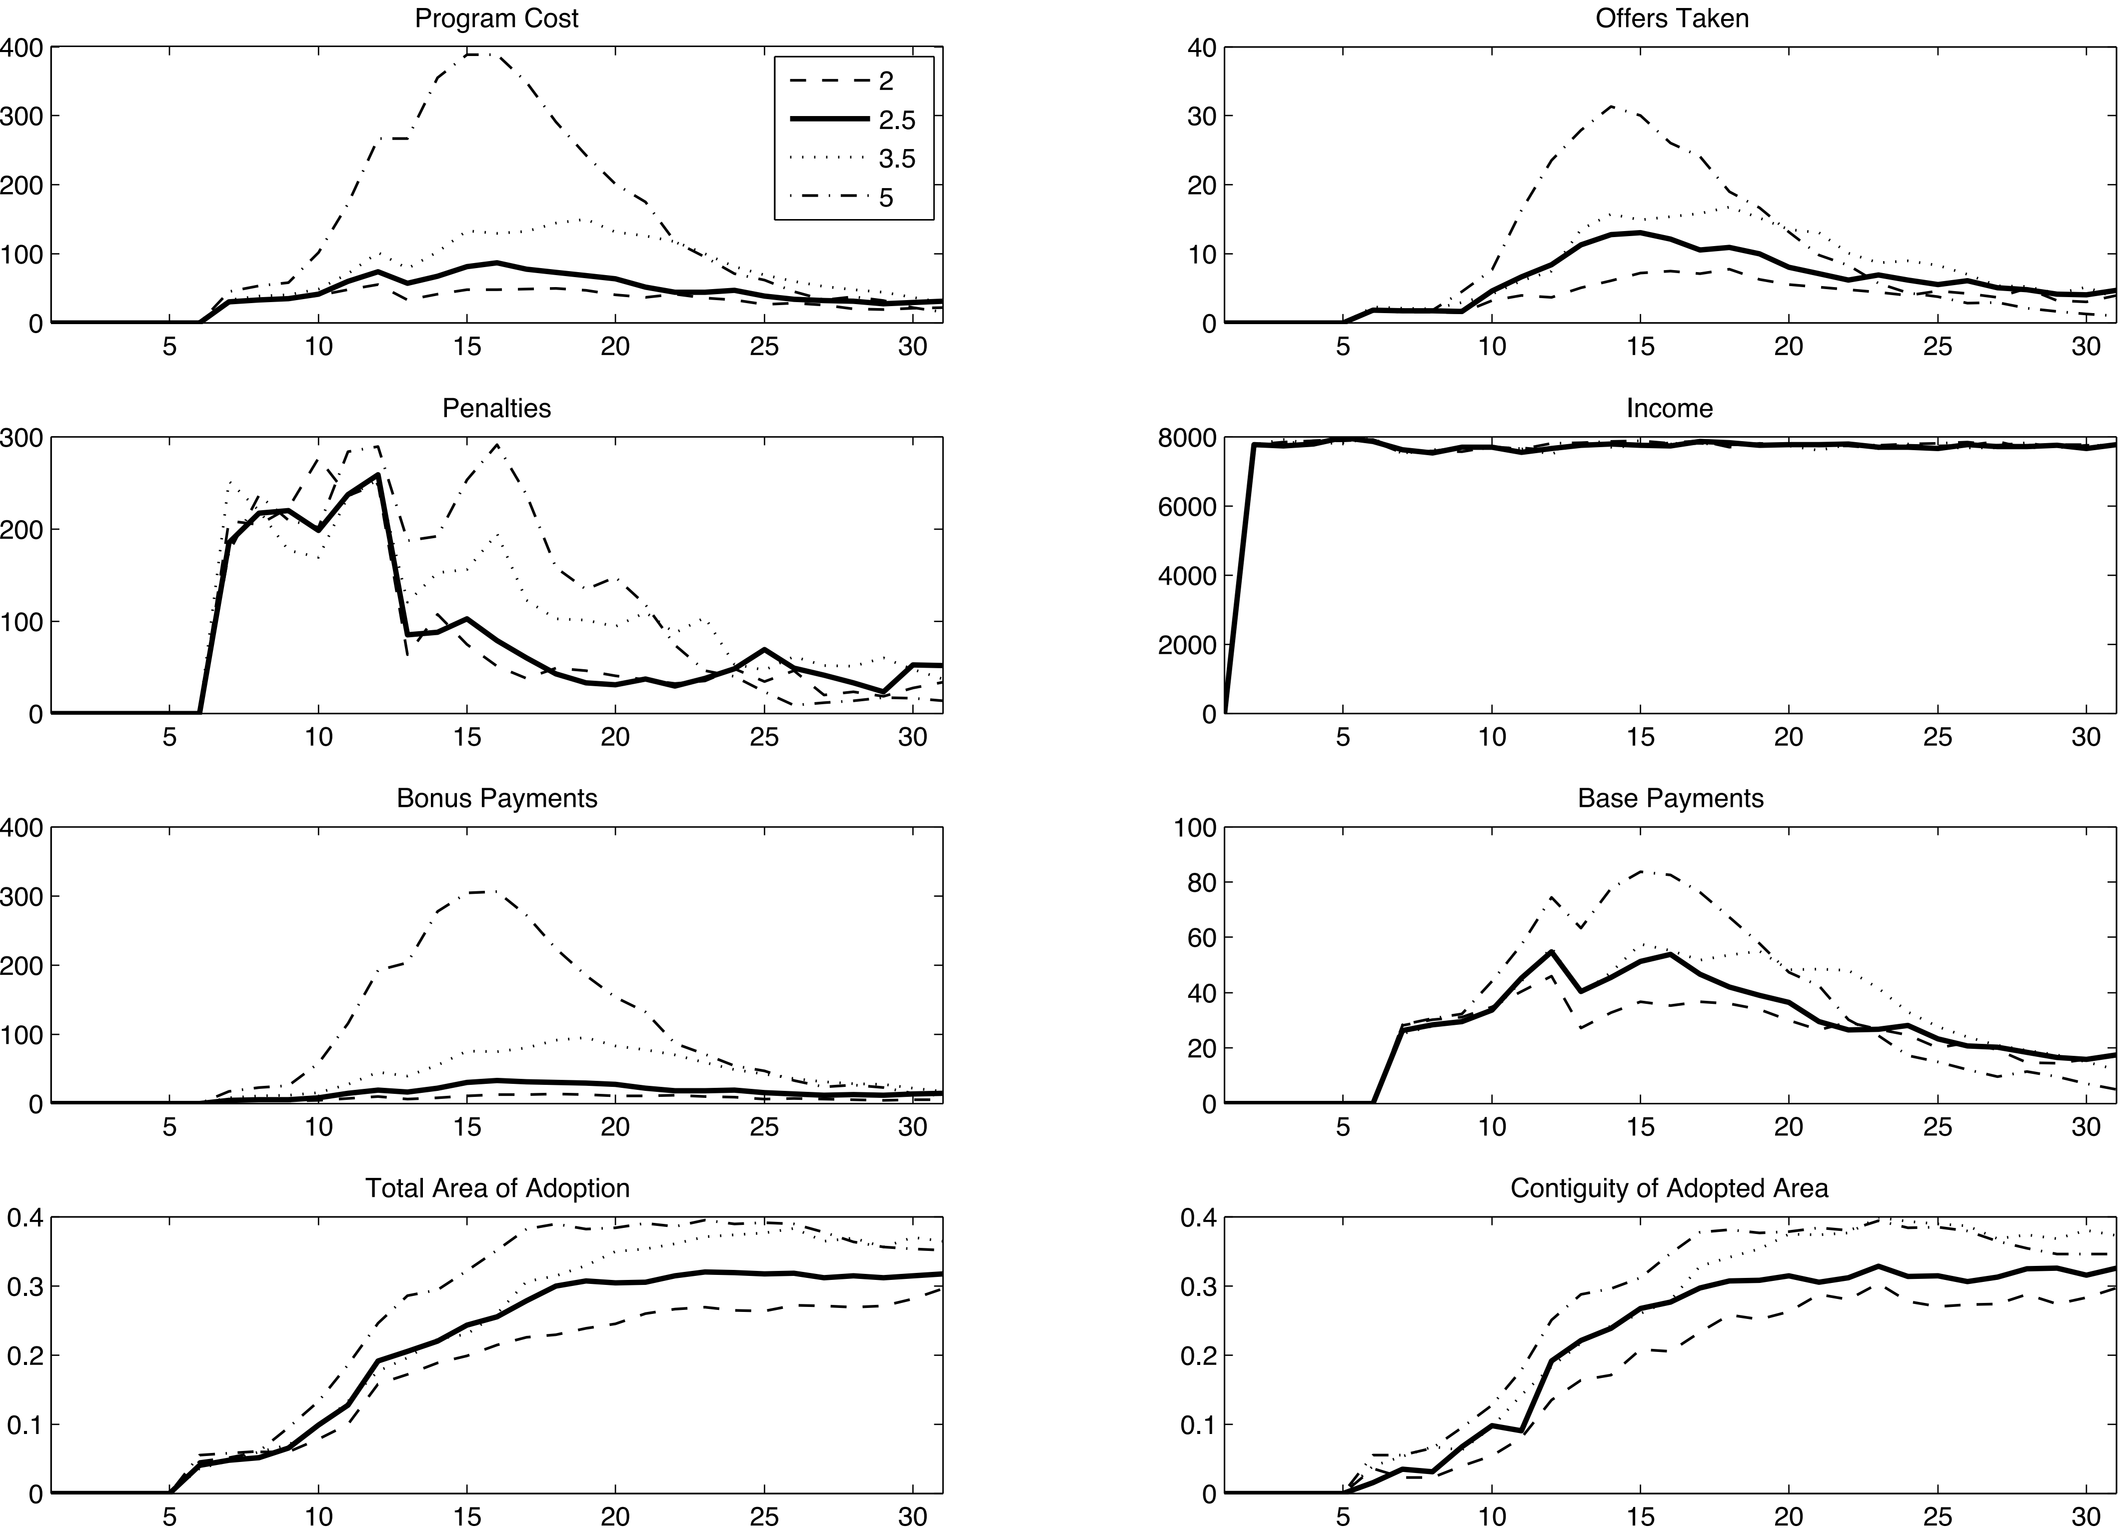

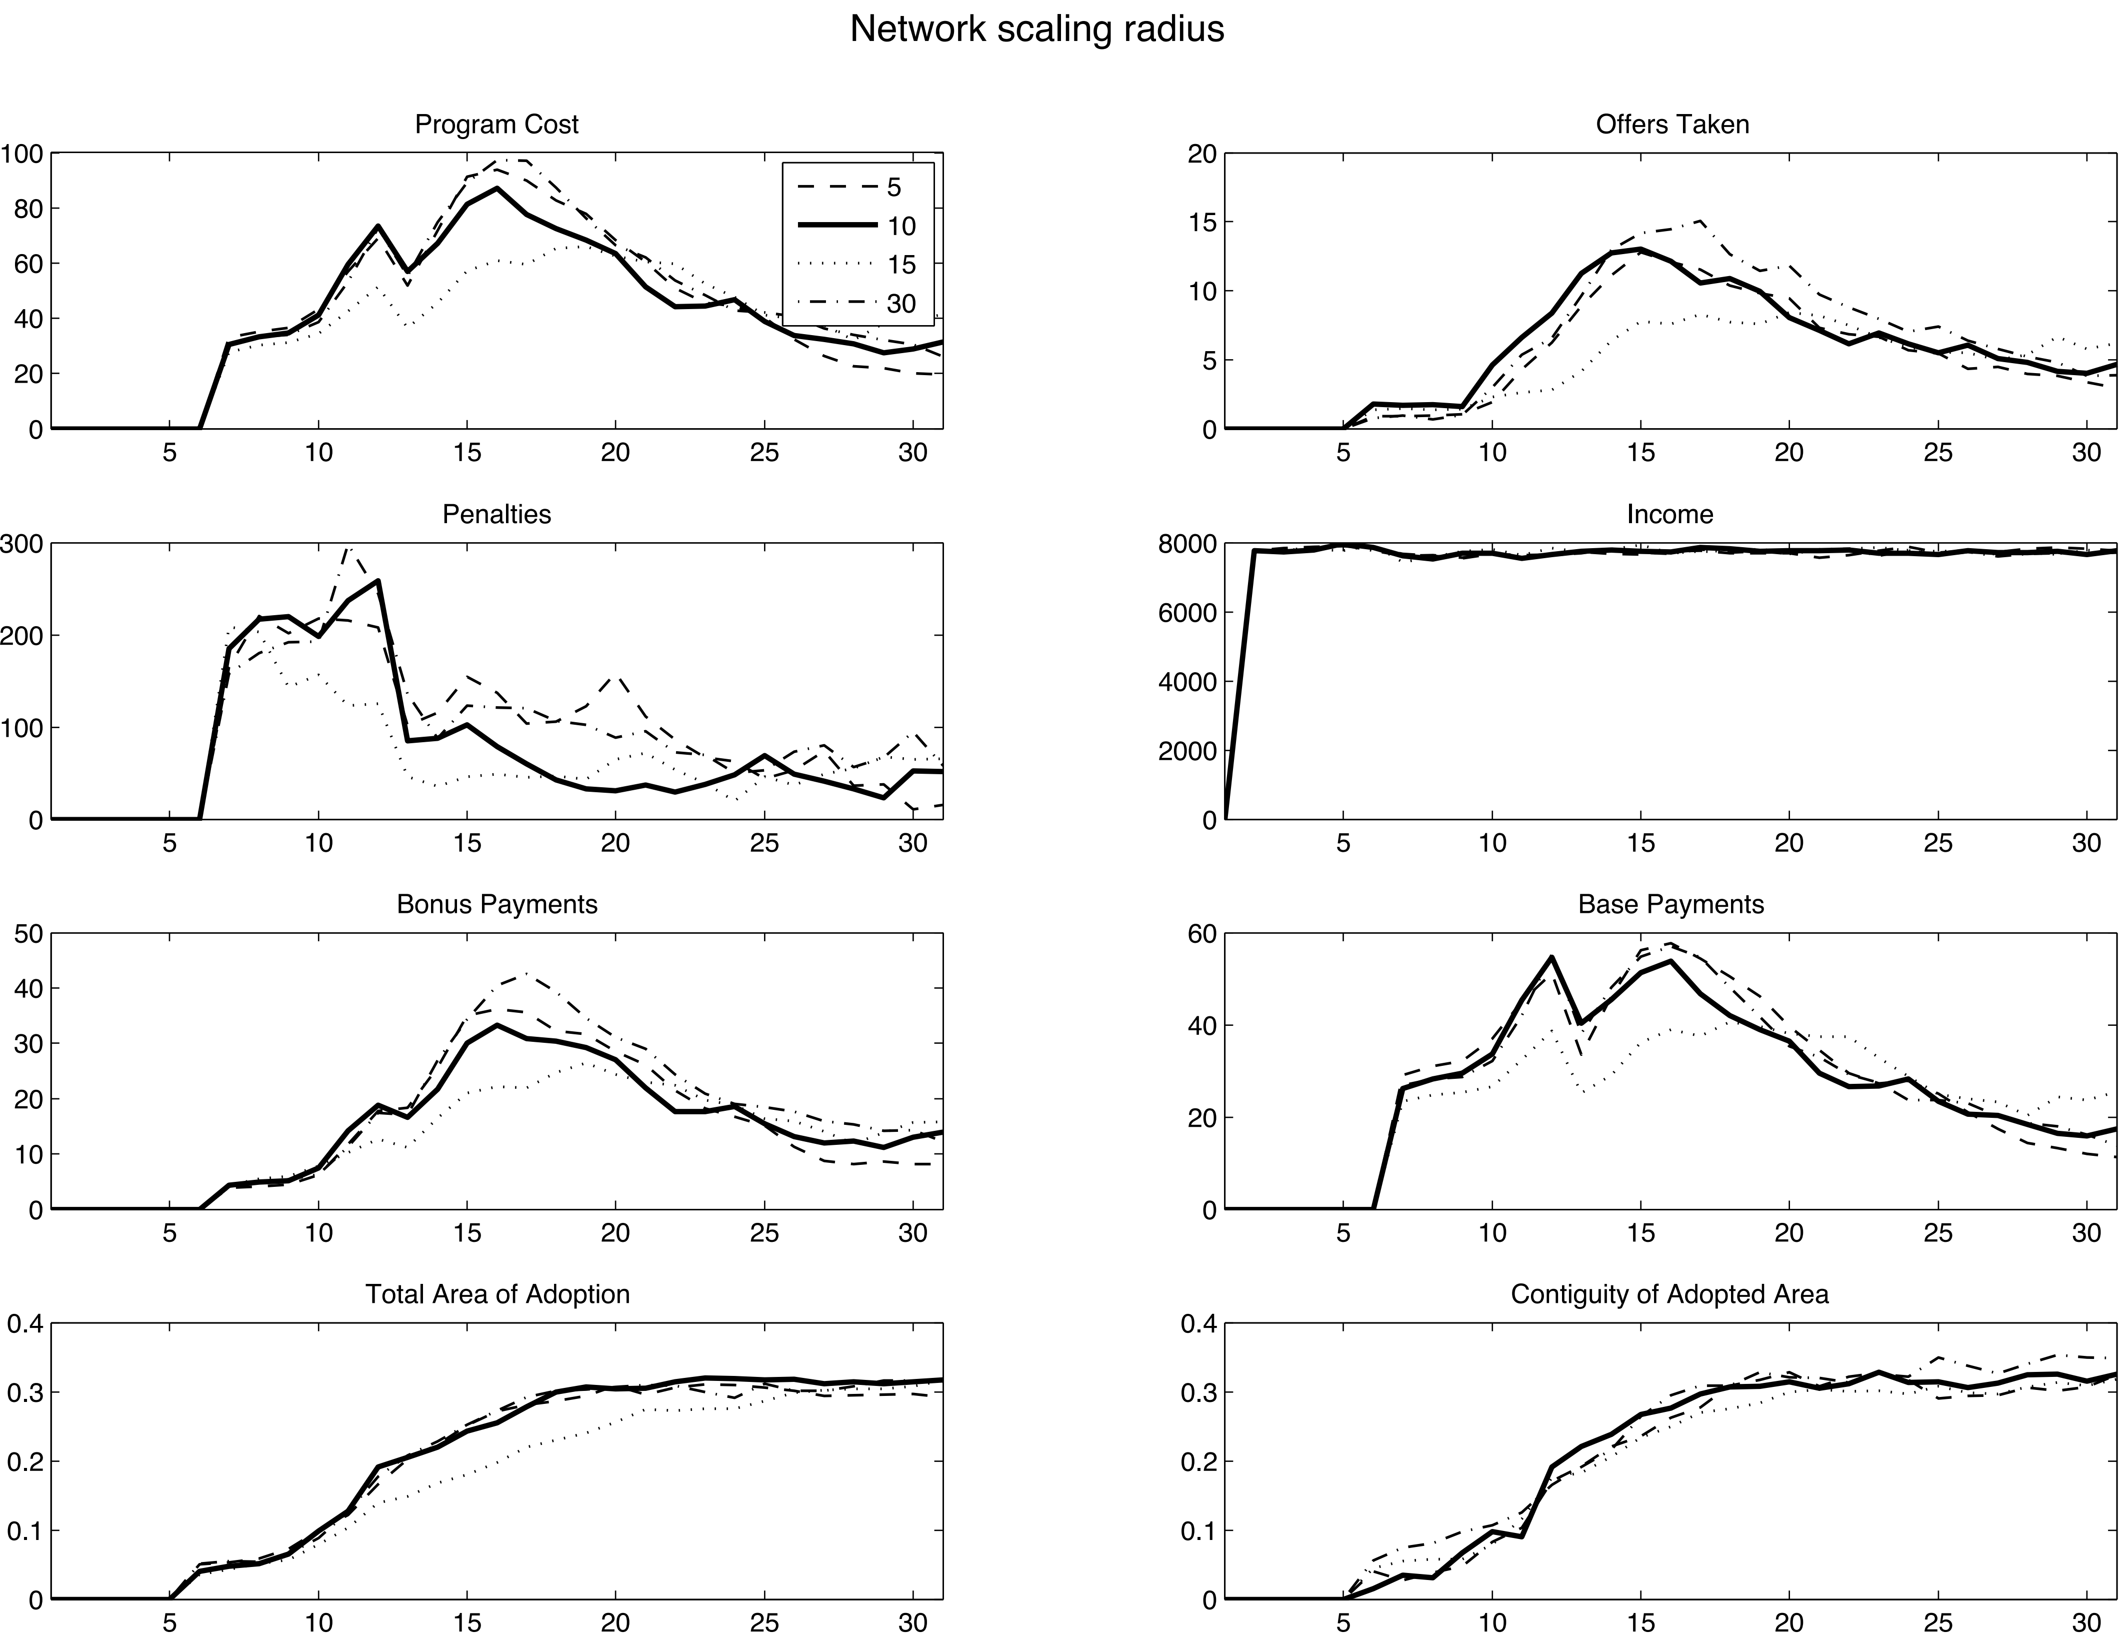

Network scaling radius
Program Cost
5
10
15
30
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

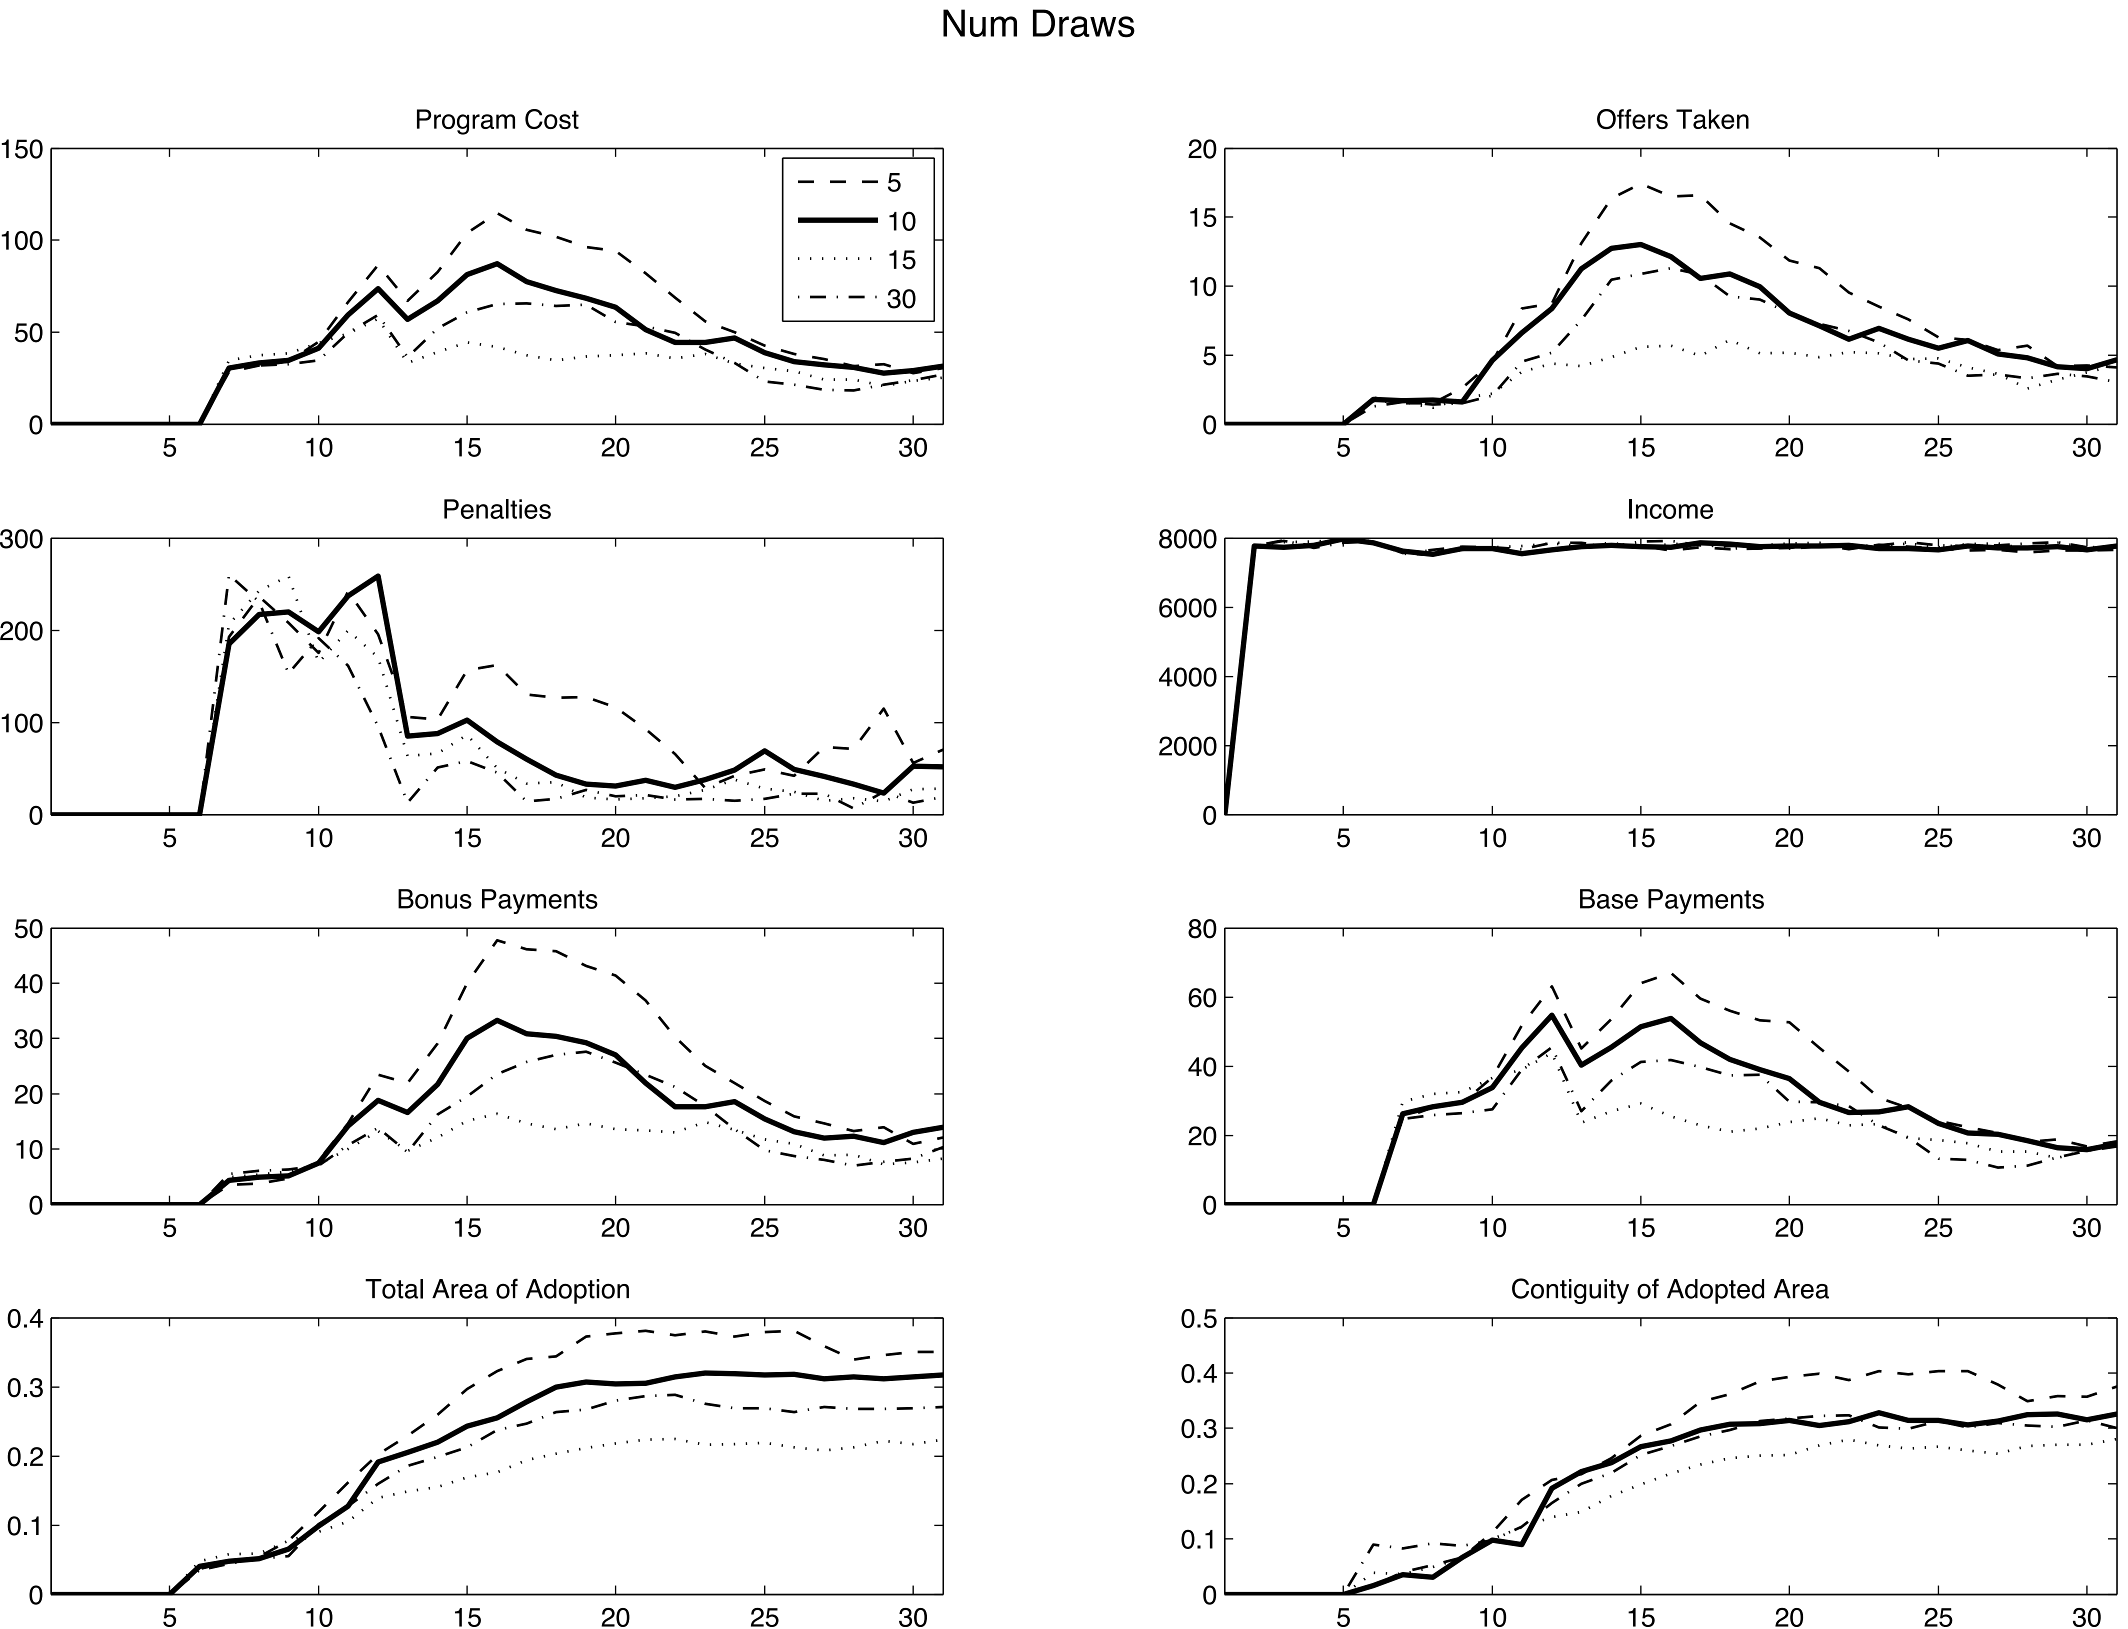
Num Draws
Program Cost
5
10
15
30
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

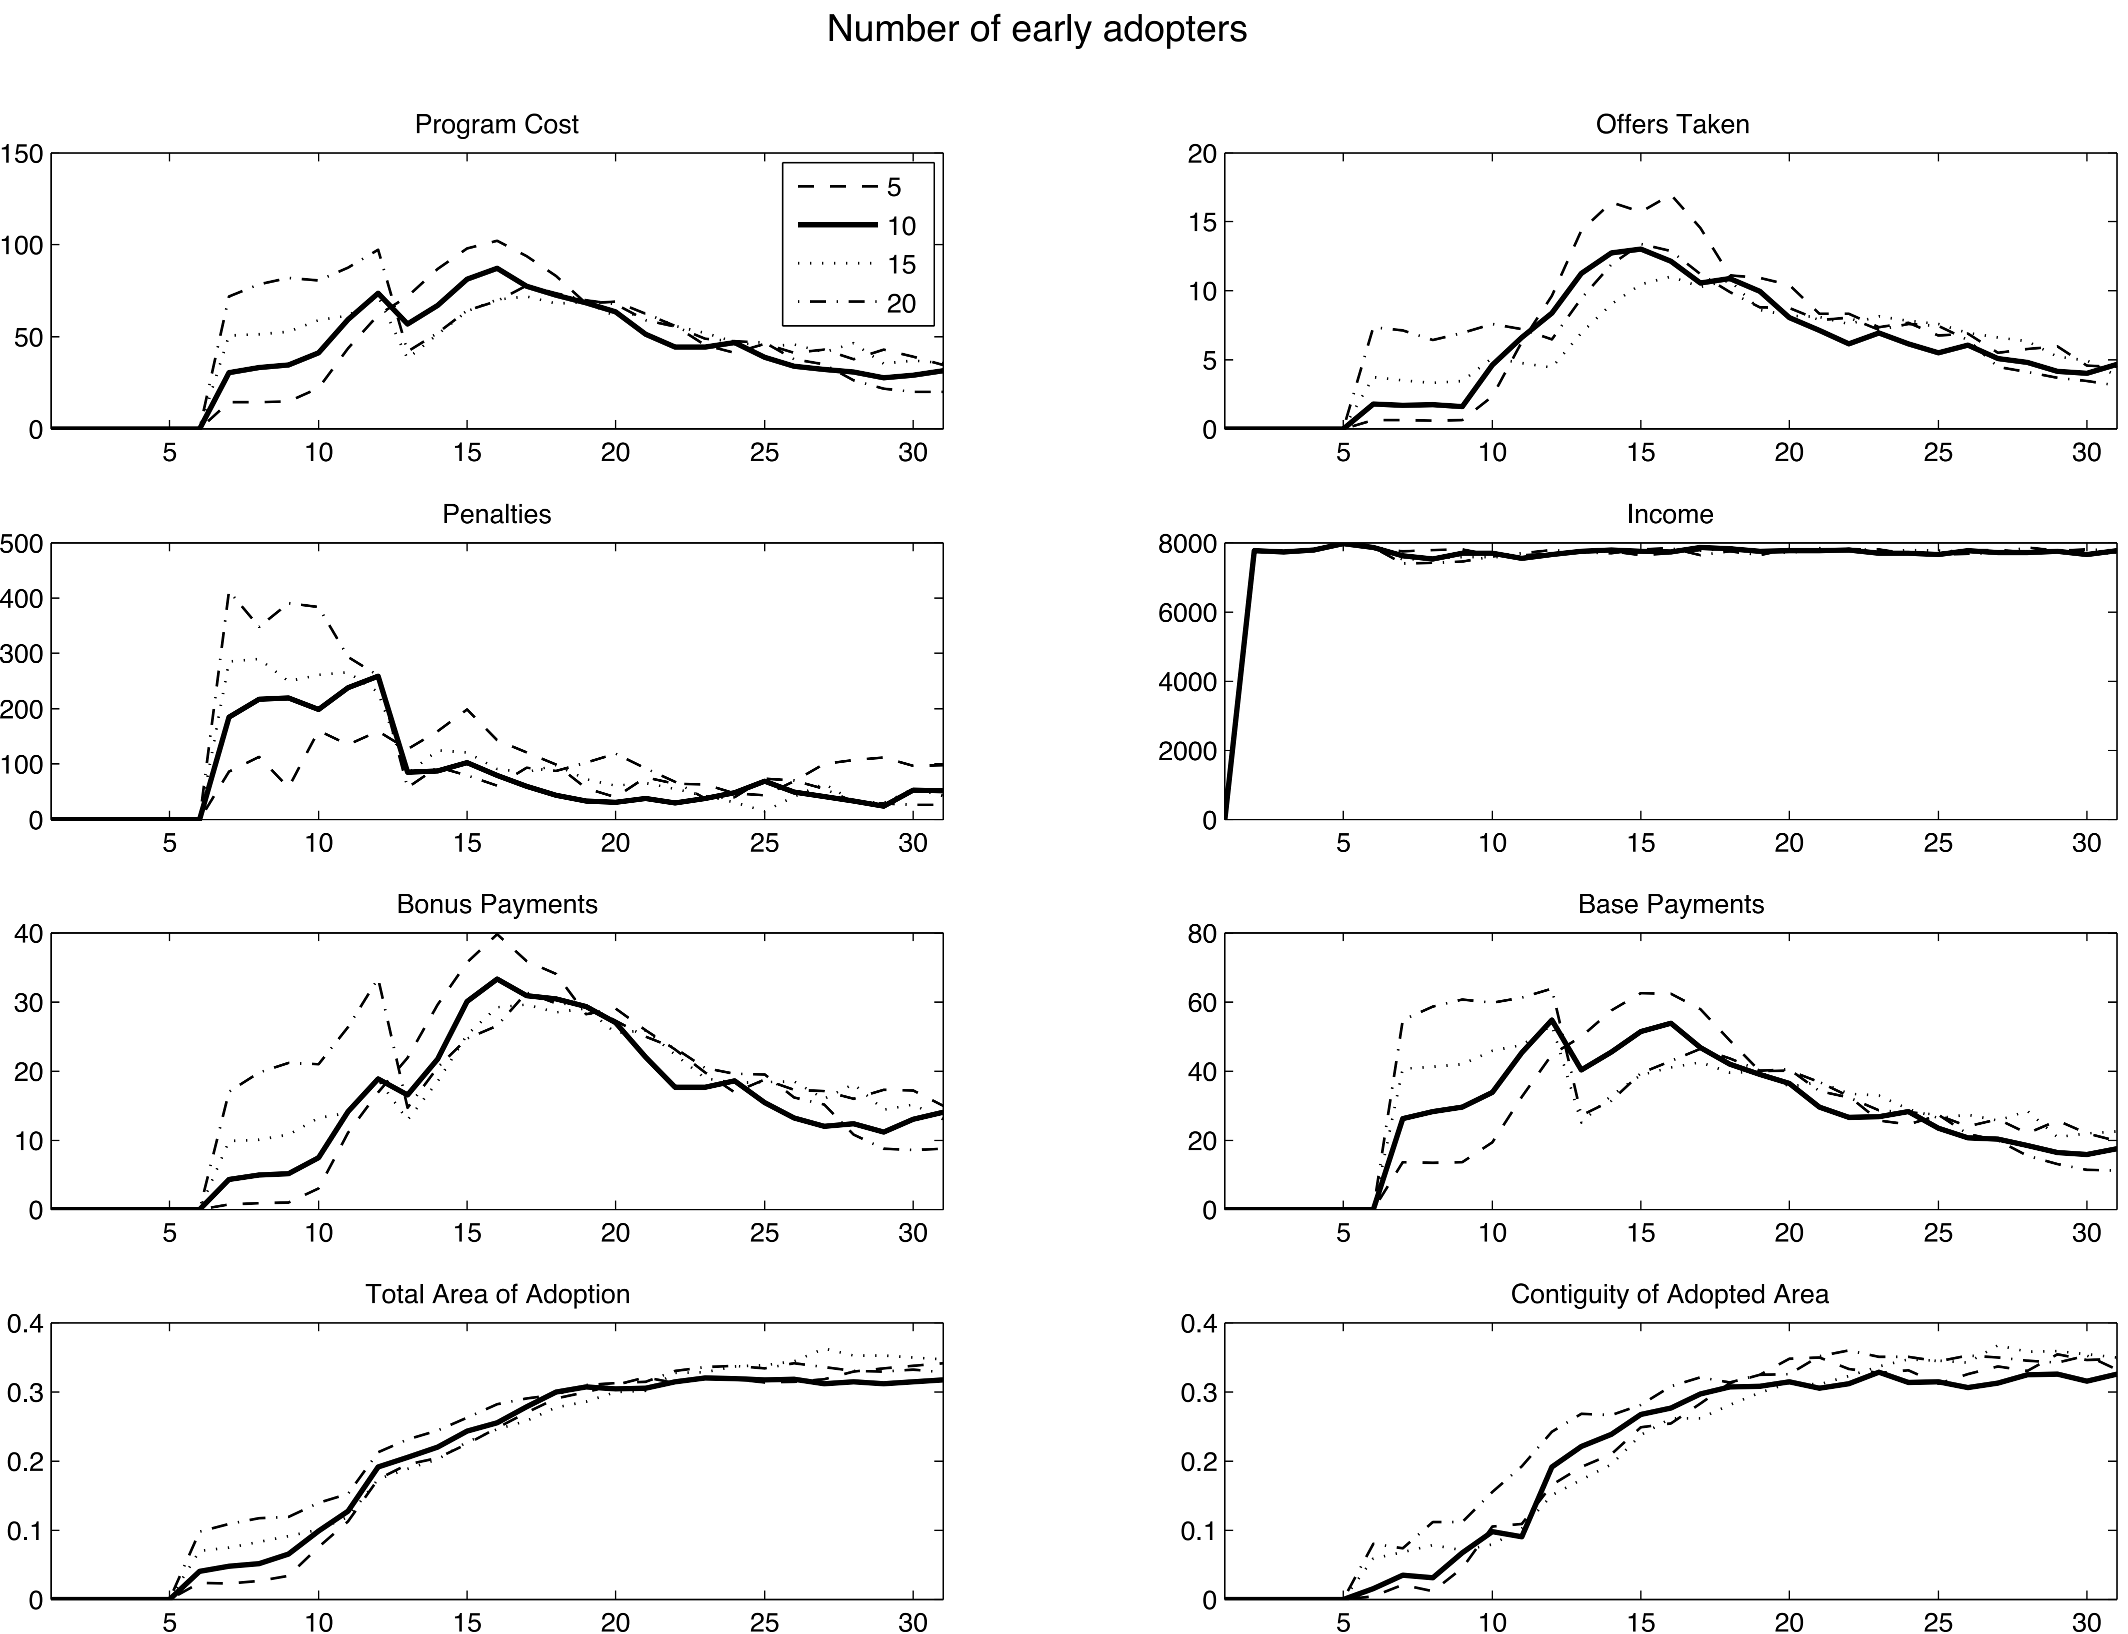
Number of early adopters
Program Cost
5
10
15
20
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

Payment Duration

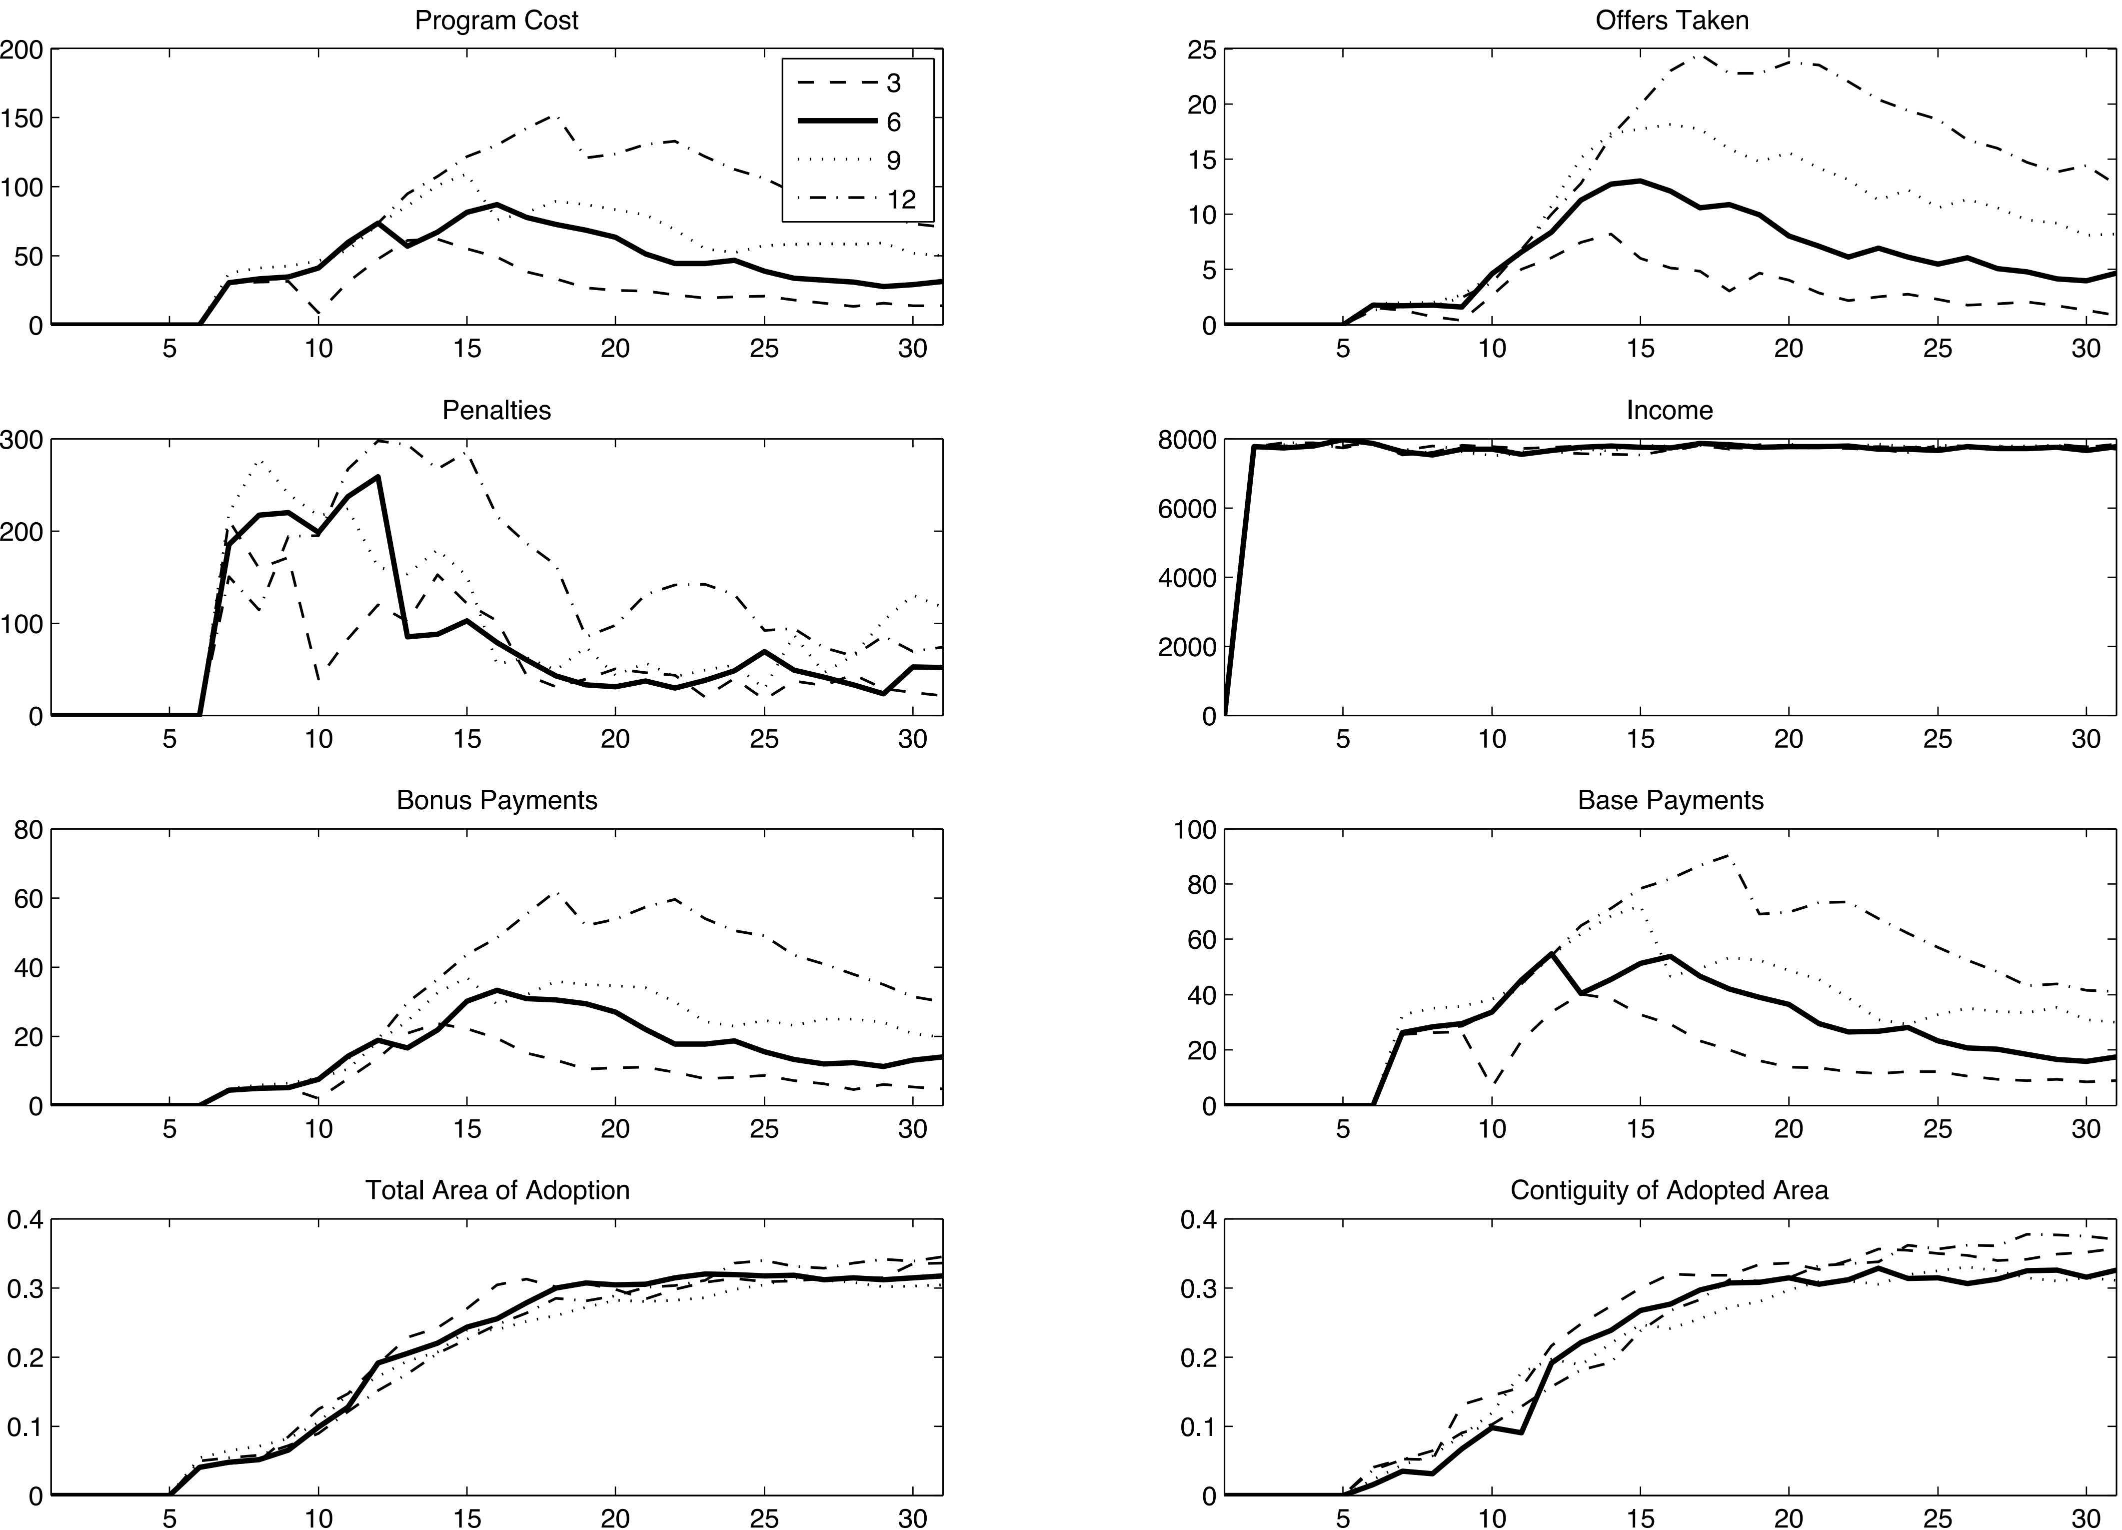

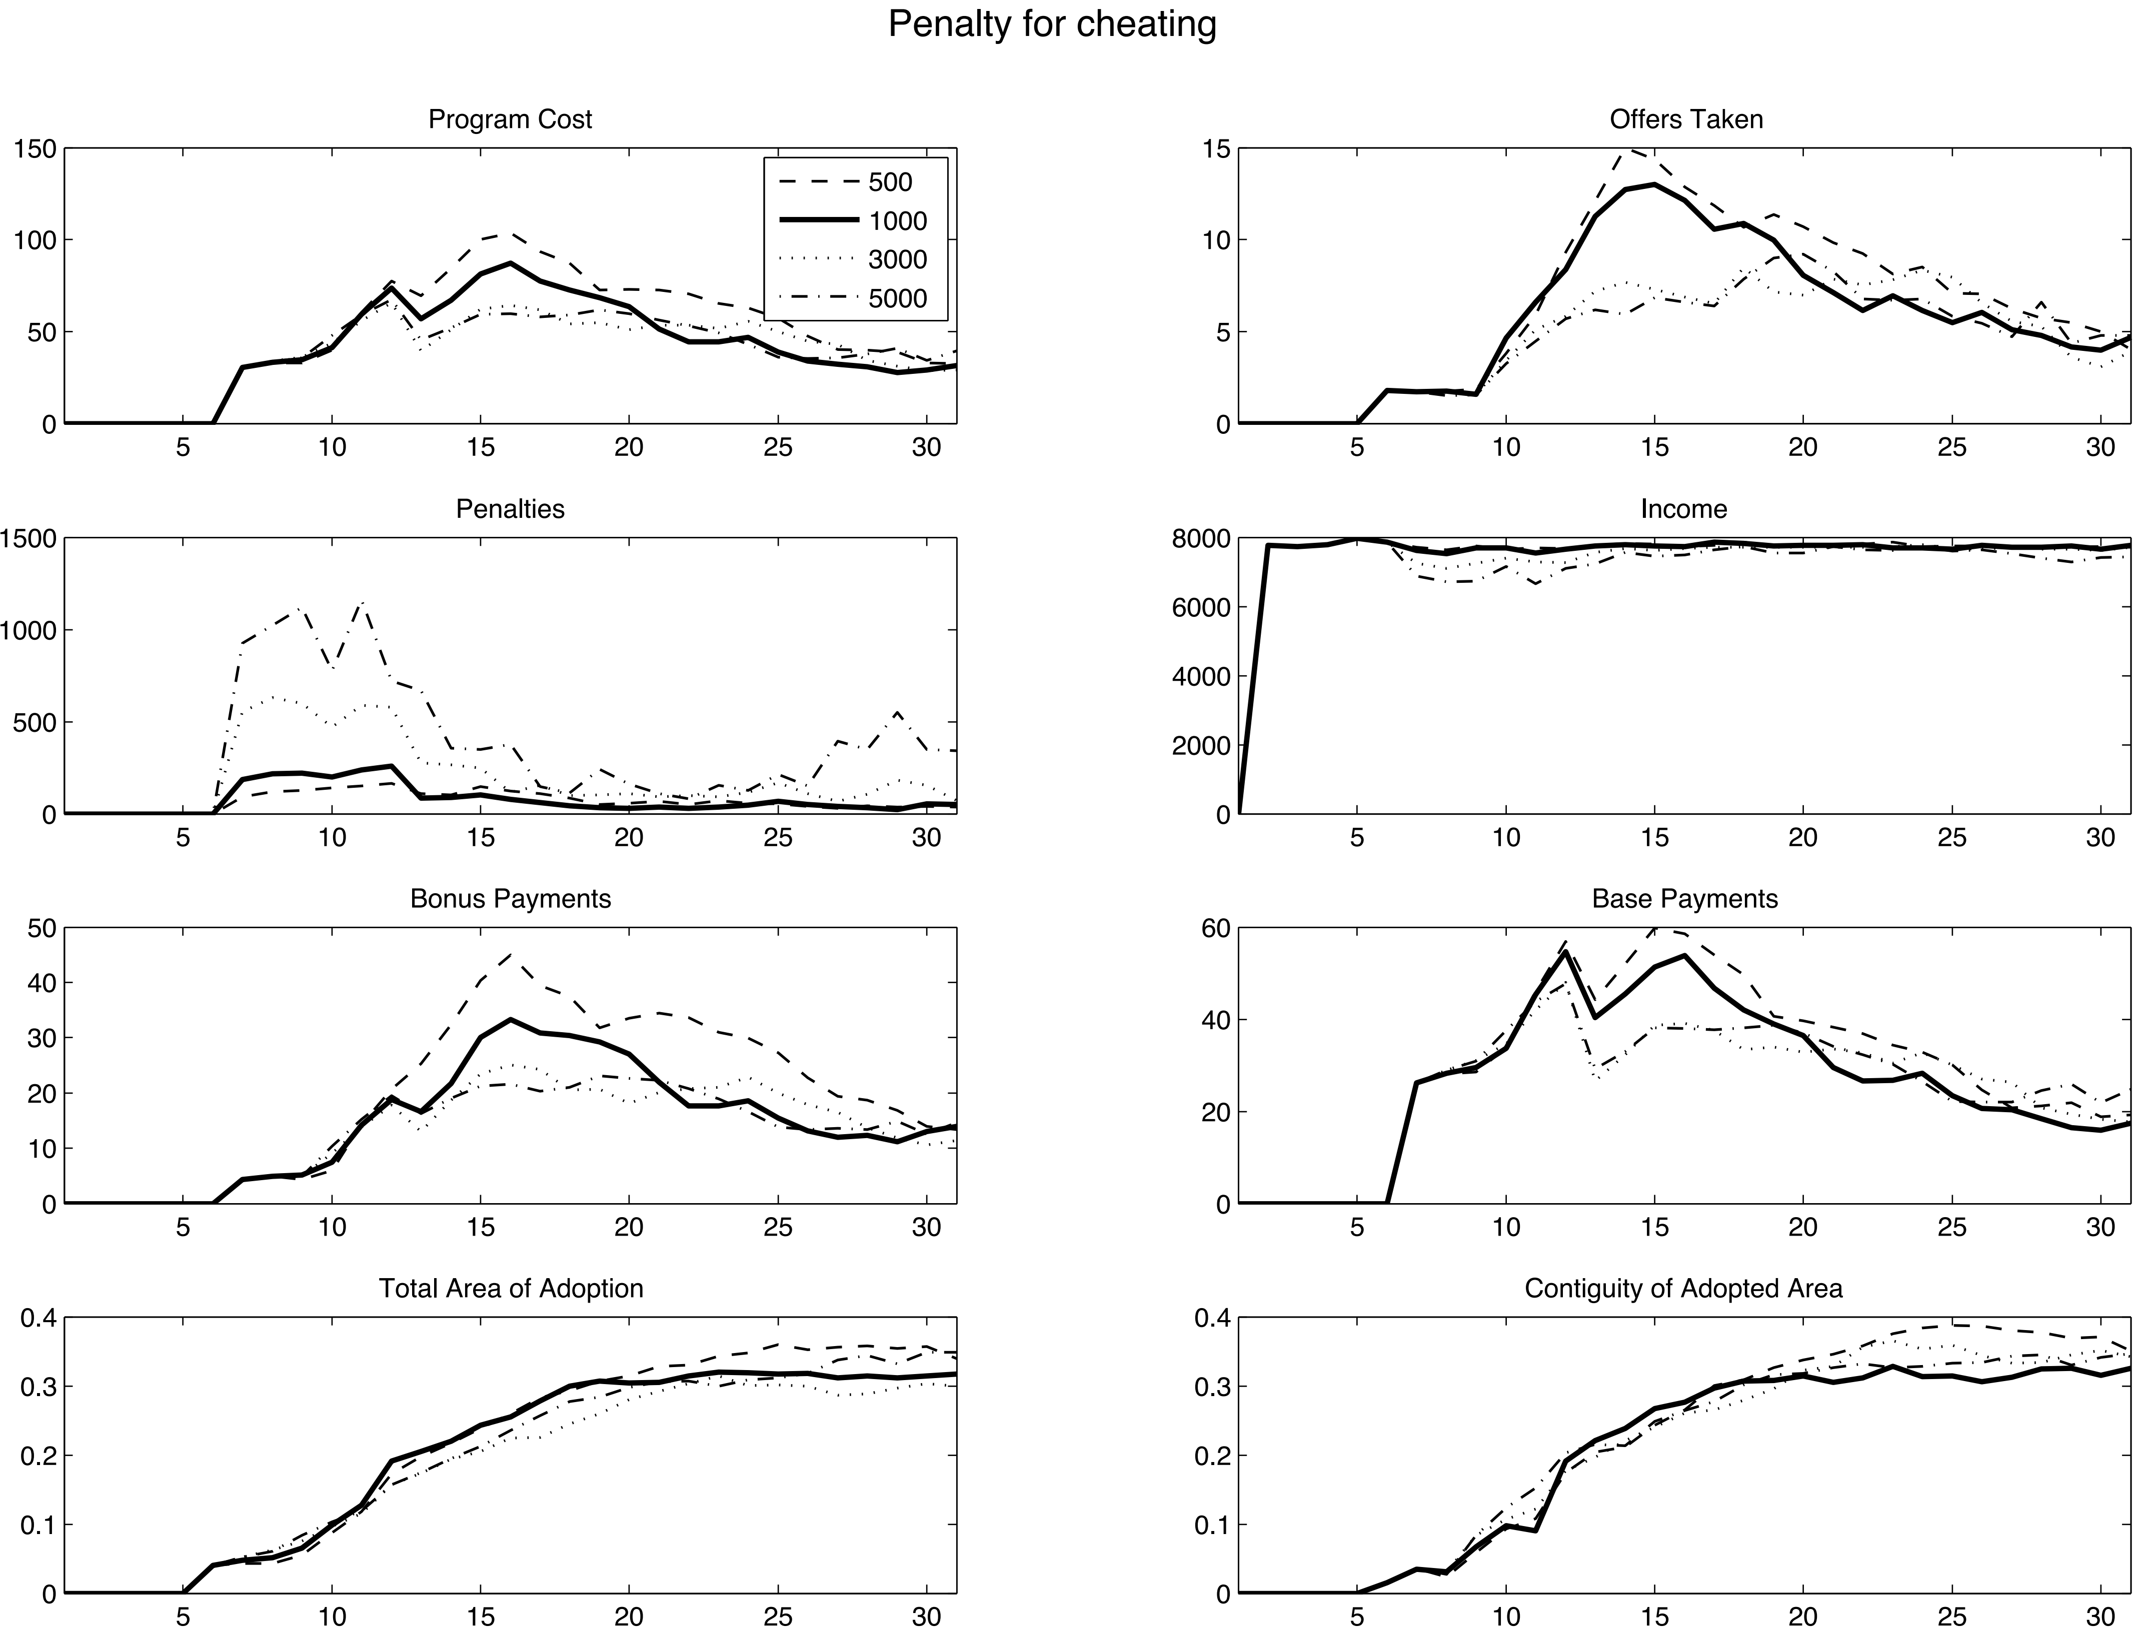
Penalty for cheating
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
500
1000
3000
5000

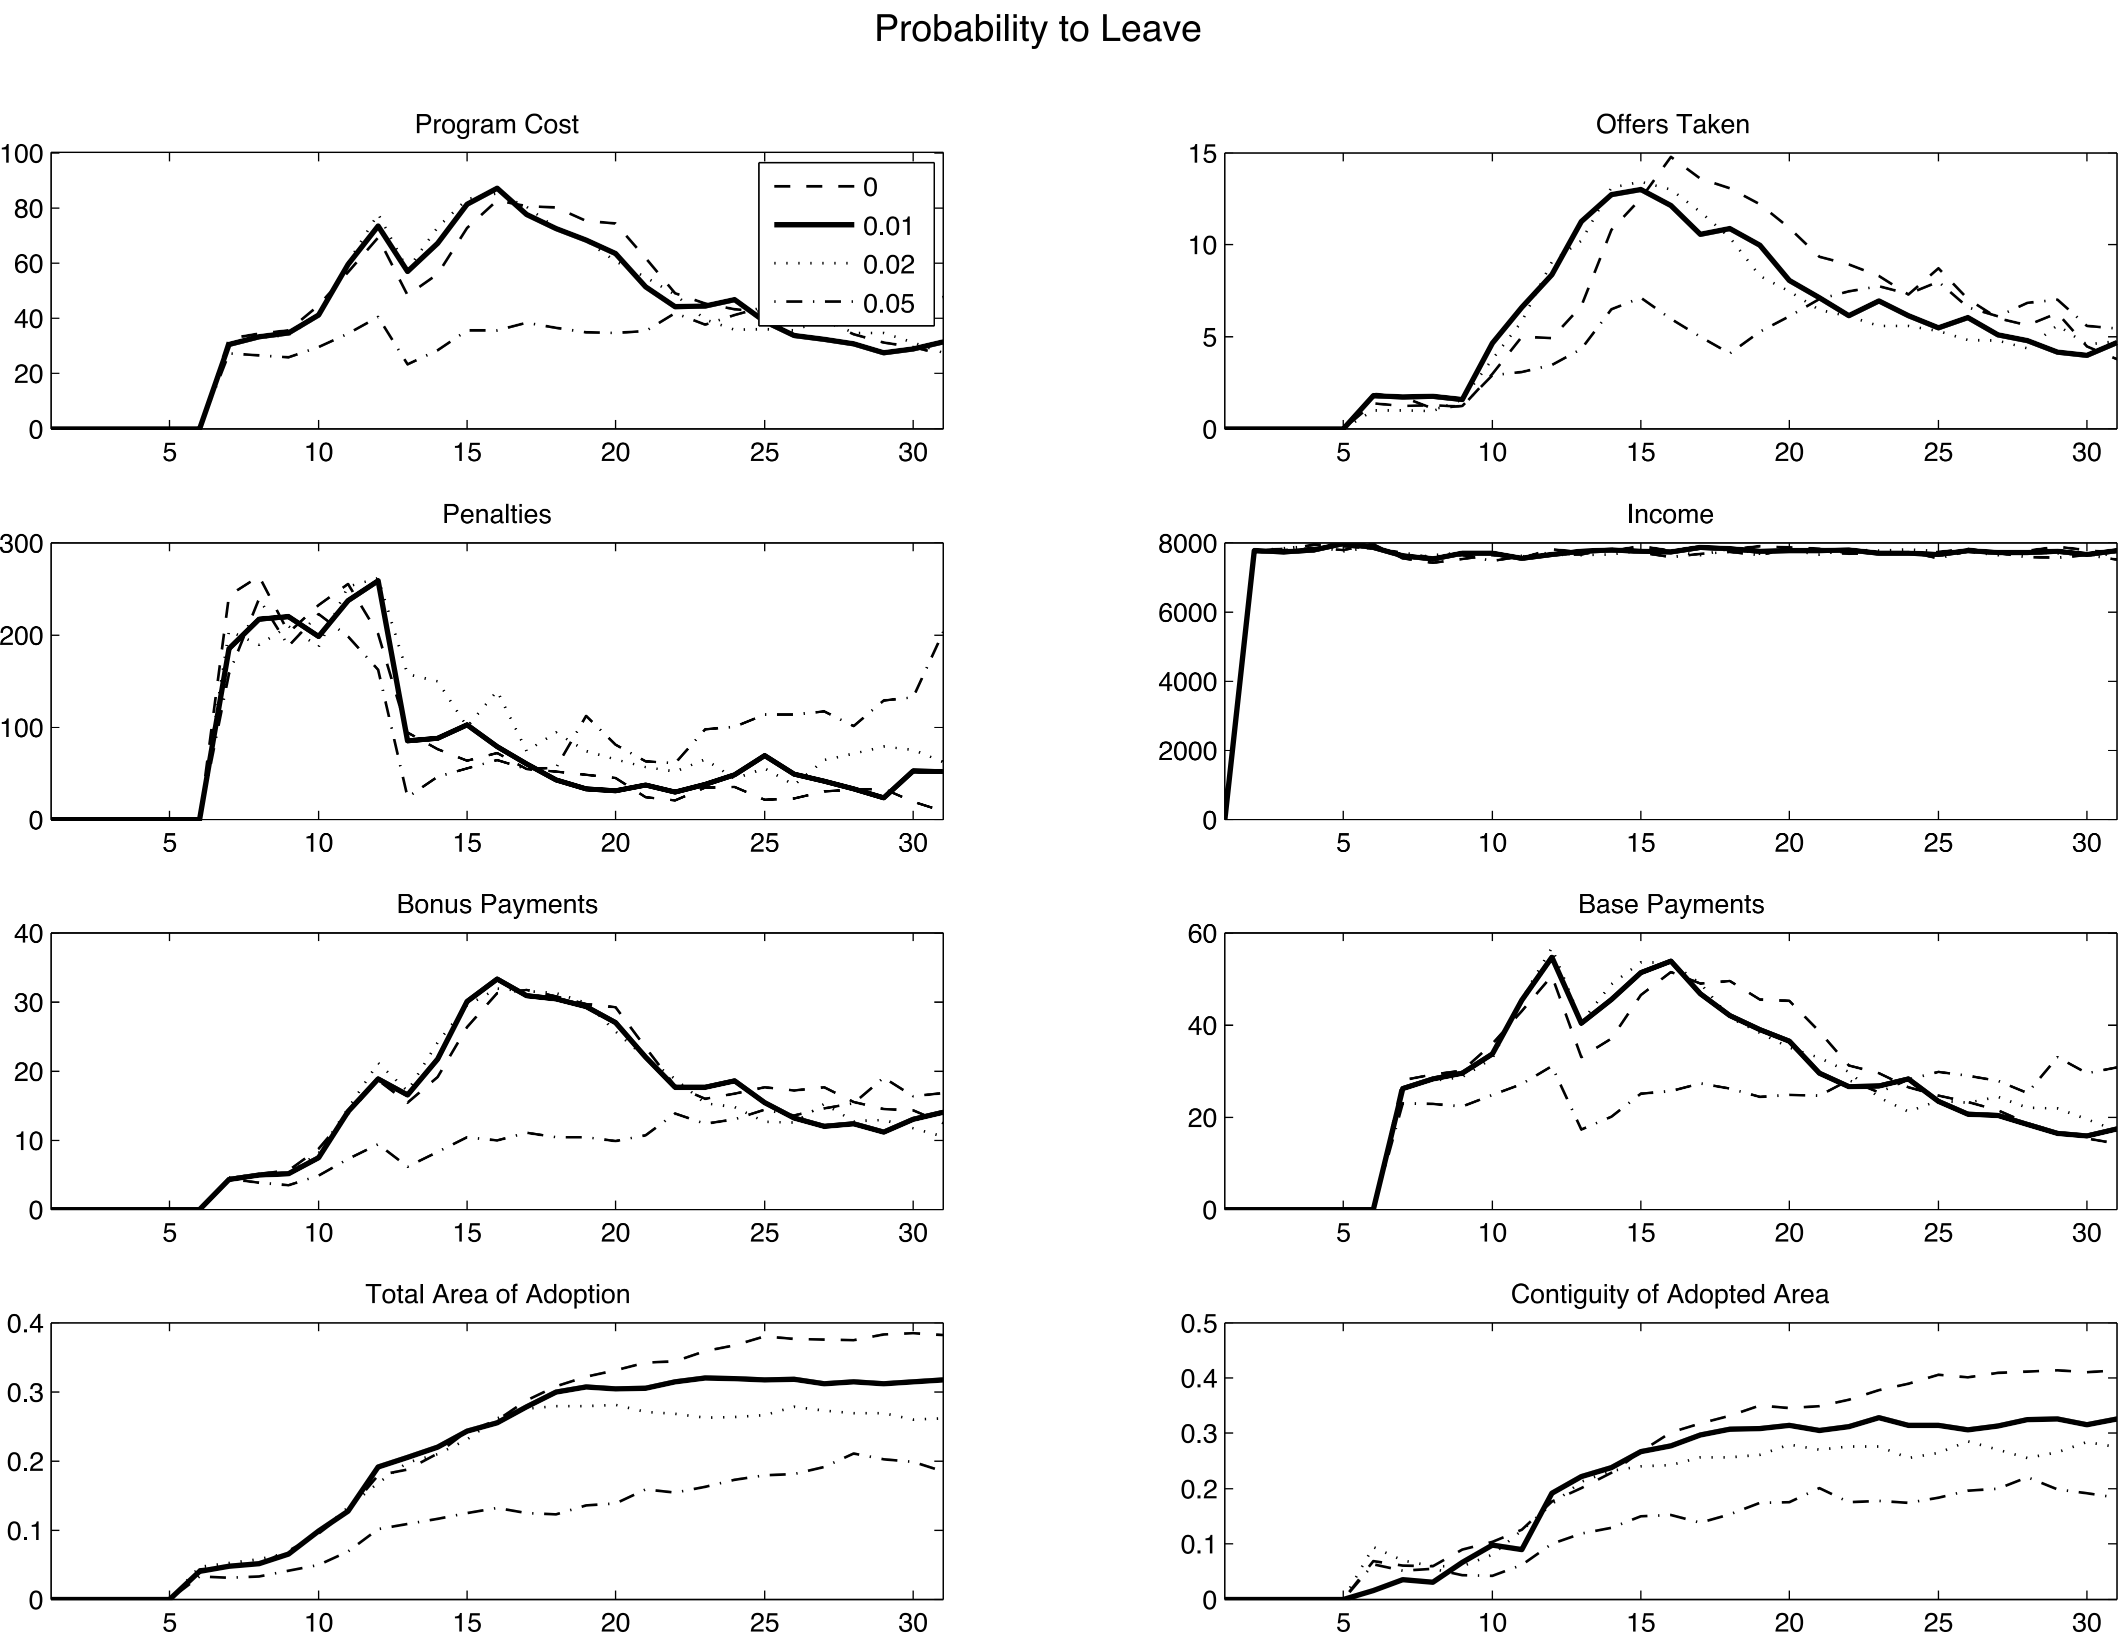

Probability to Leave
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
0
0.01
0.02
0.05

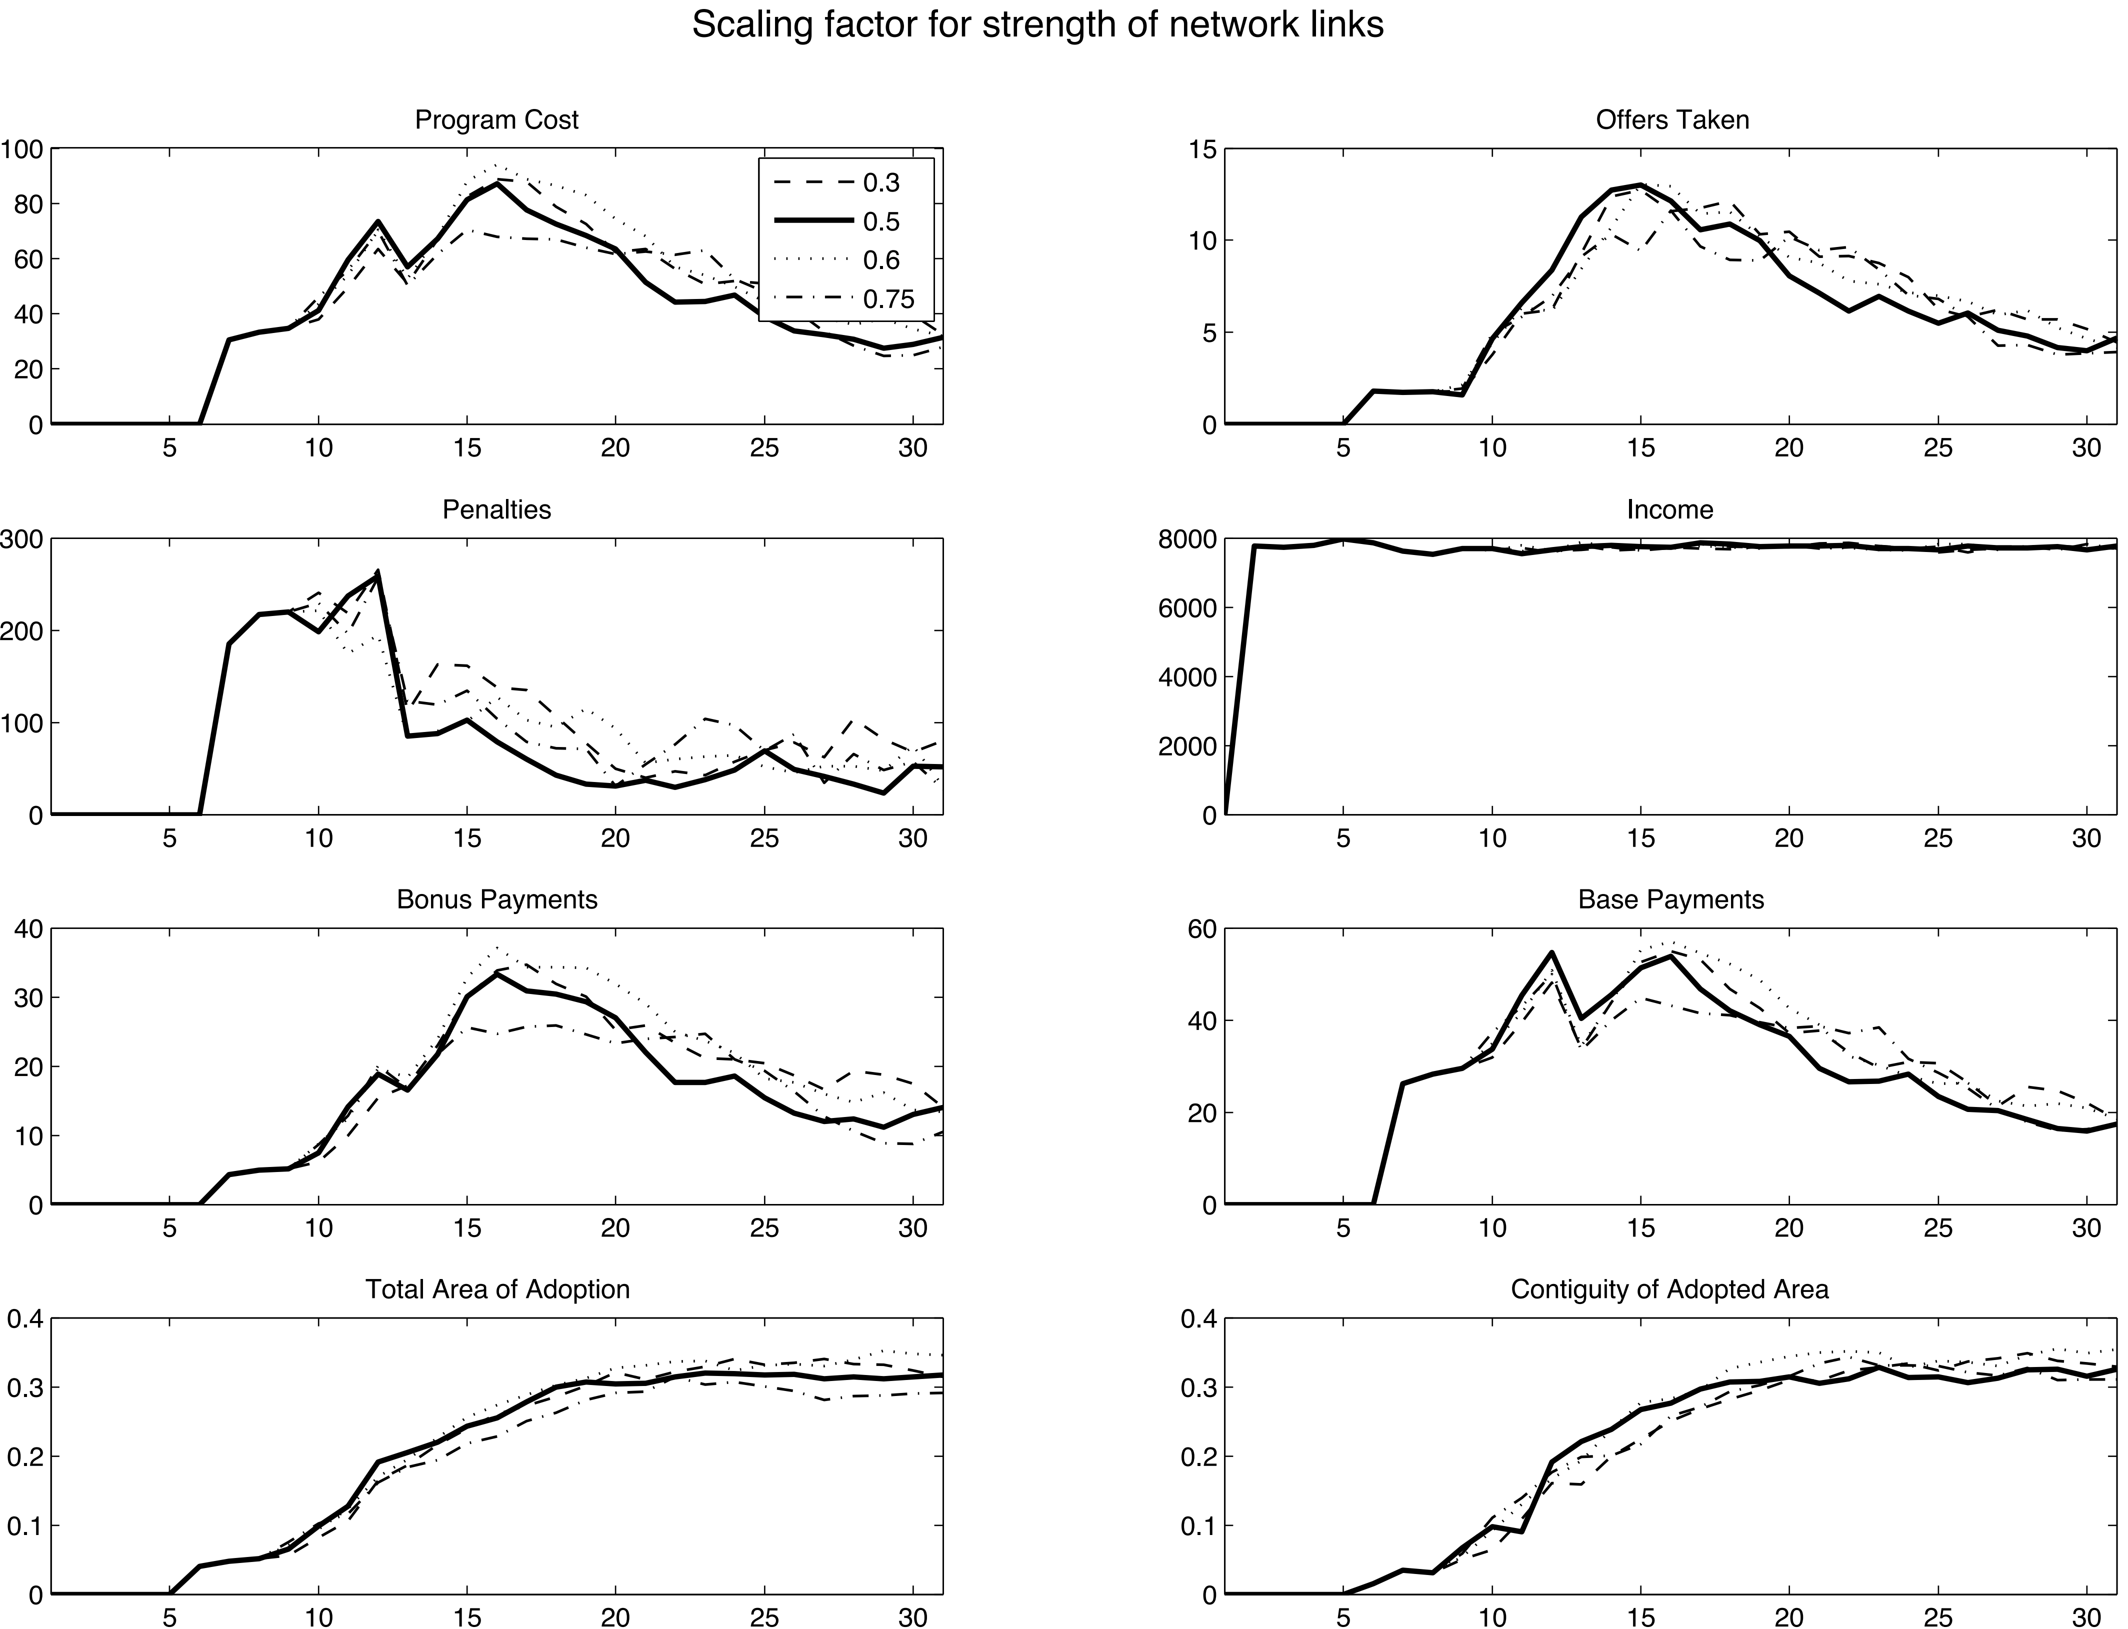
Scaling factor for strength of network links
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
0.3
0.5
0.6
0.75

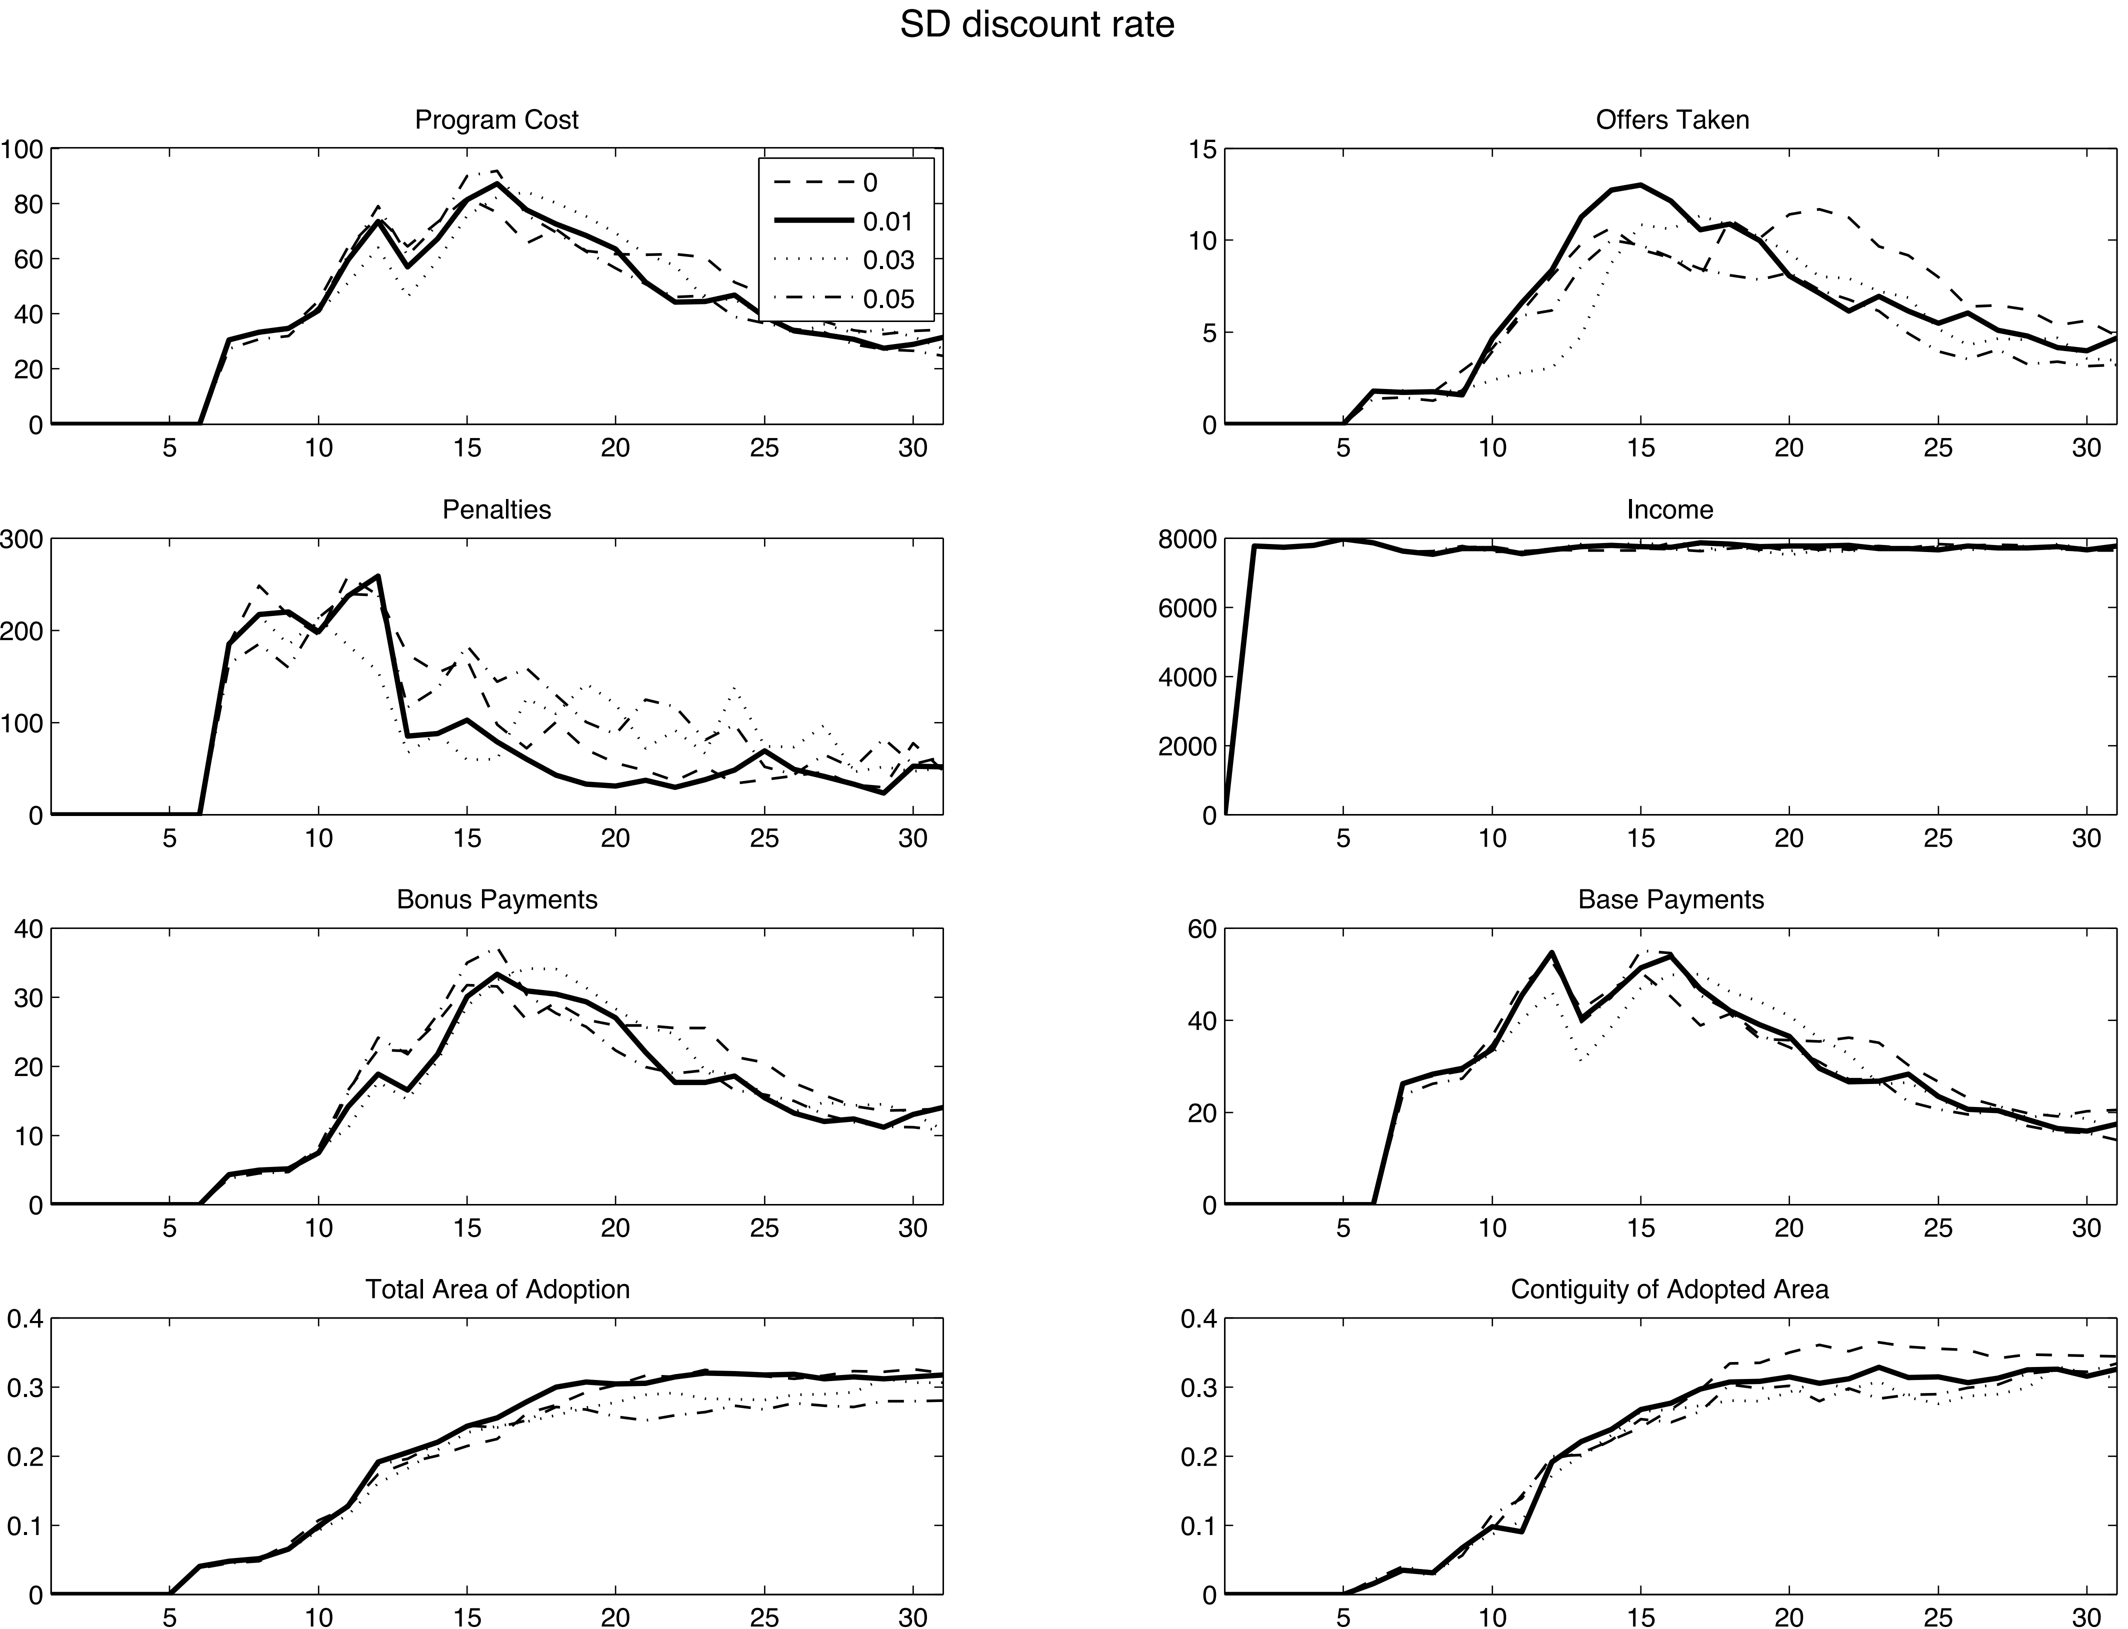
SD discount rate
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
0
0.01
0.03
0.05

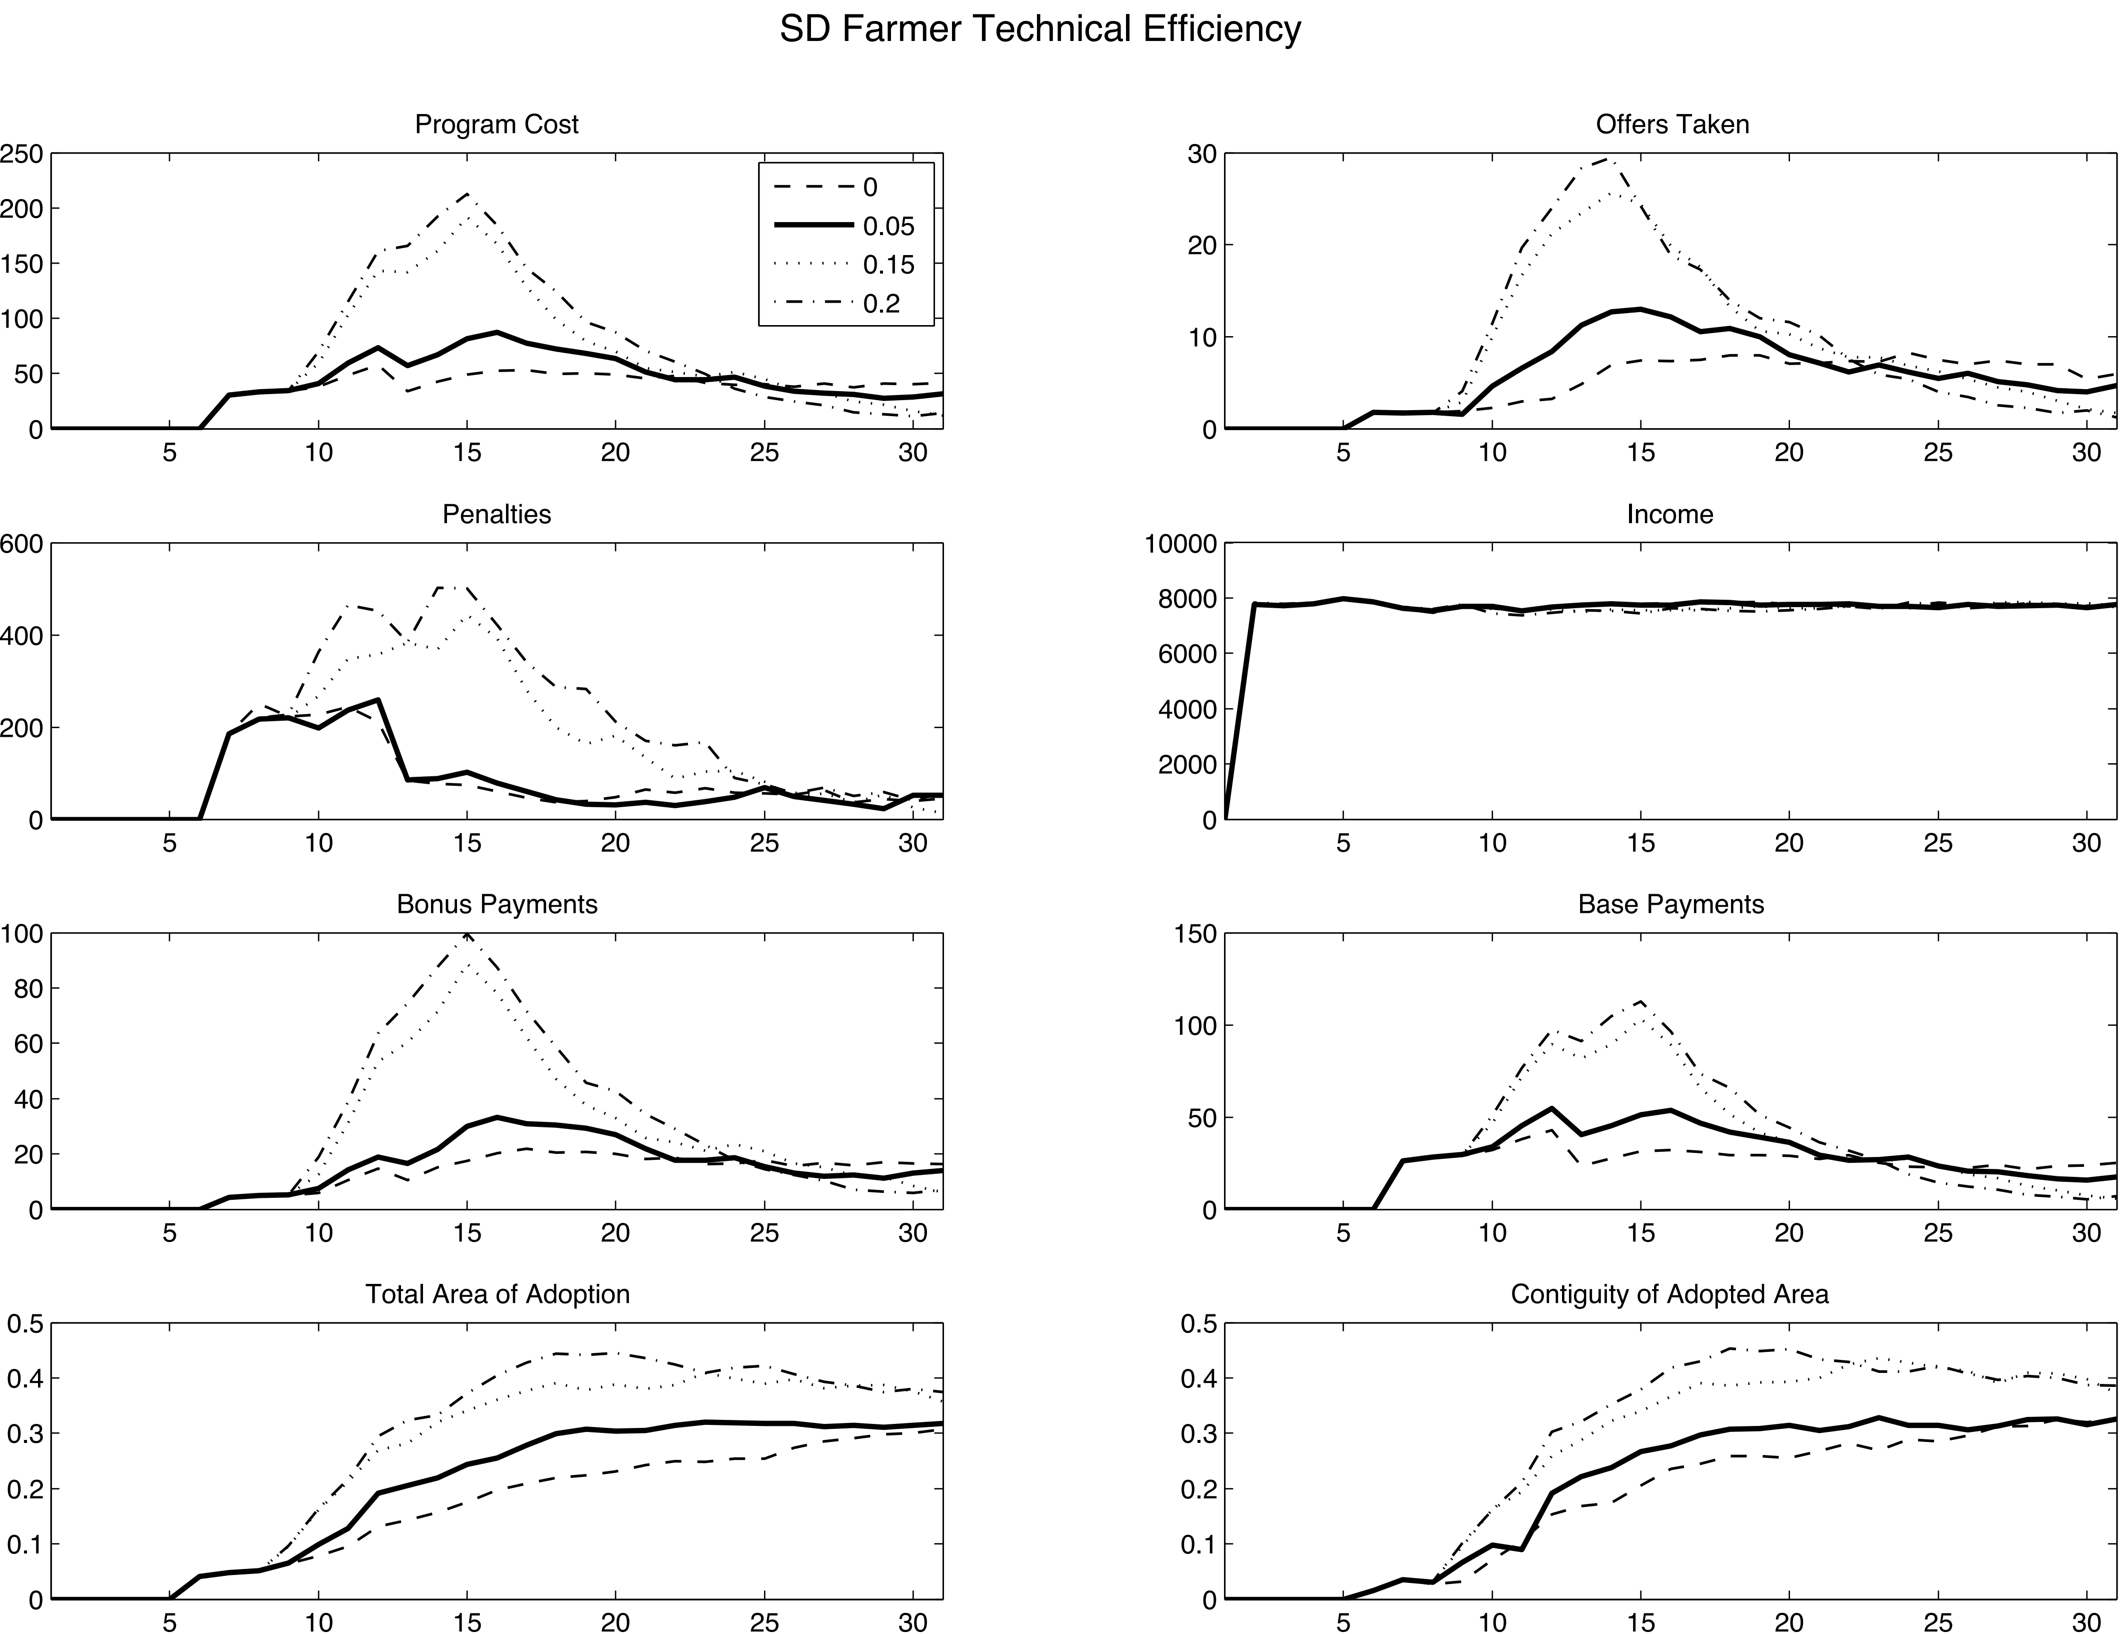

SD Farmer Technical Efficiency
Program Cost
0
0.05
0.15
0.2
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

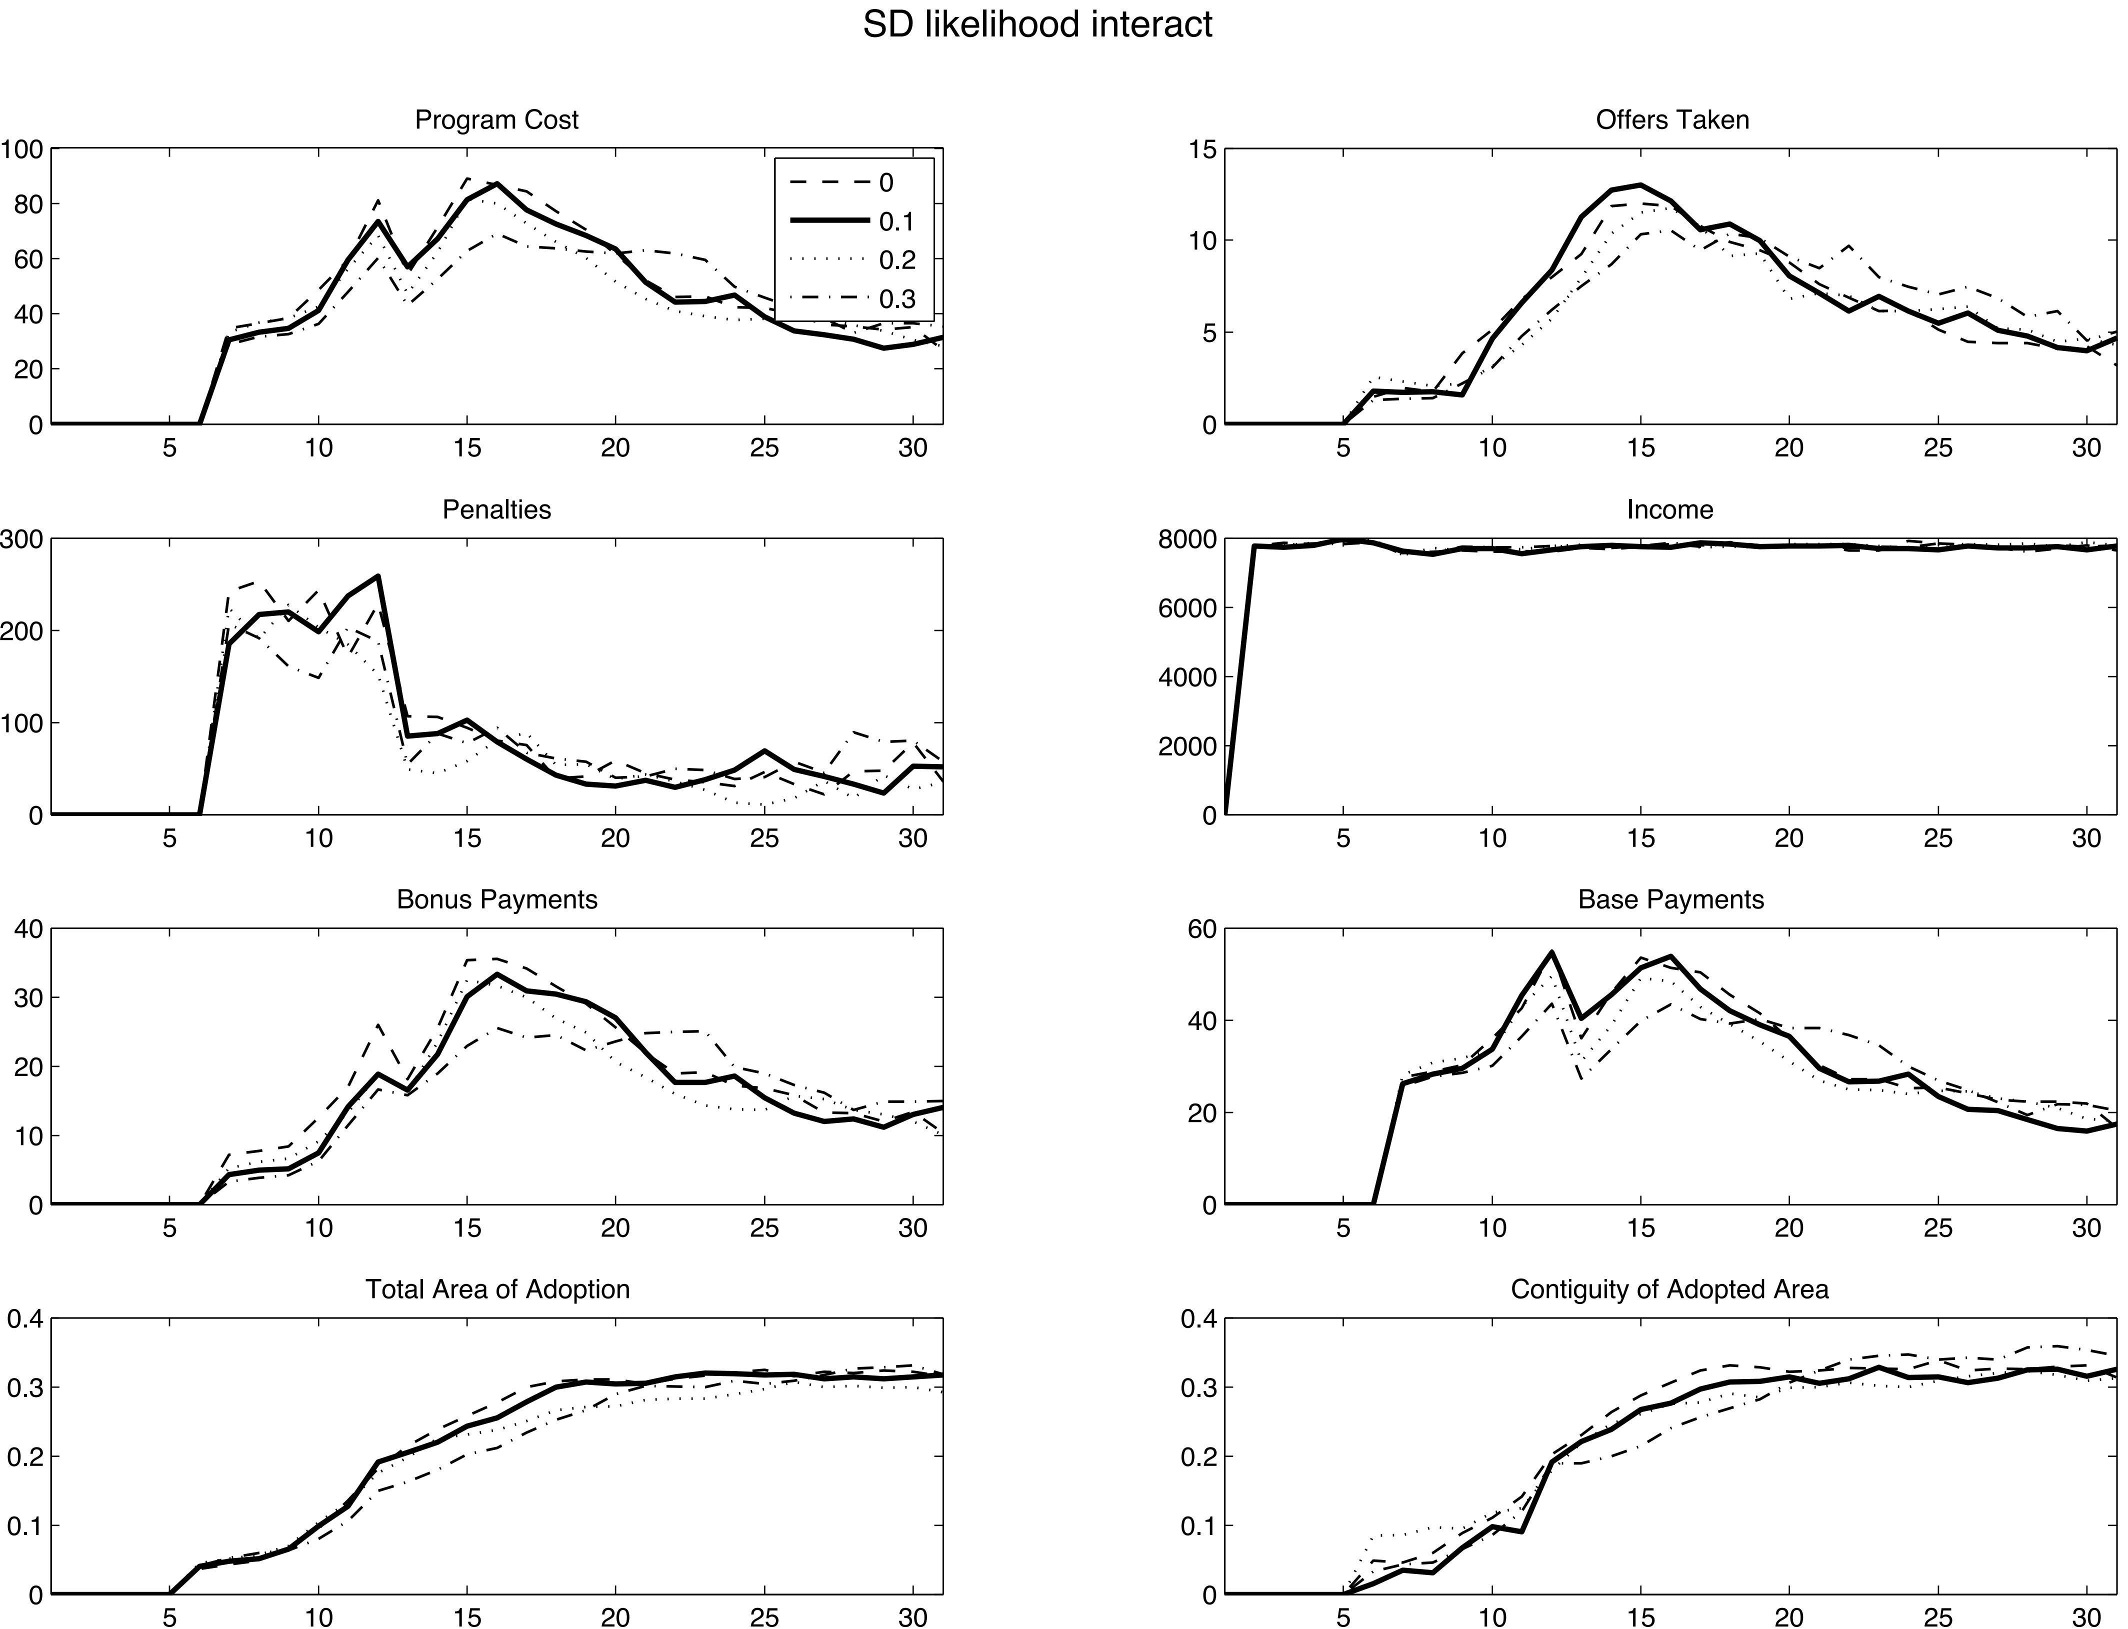

SD likelihood interact
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
0
0.1
0.2
0.3

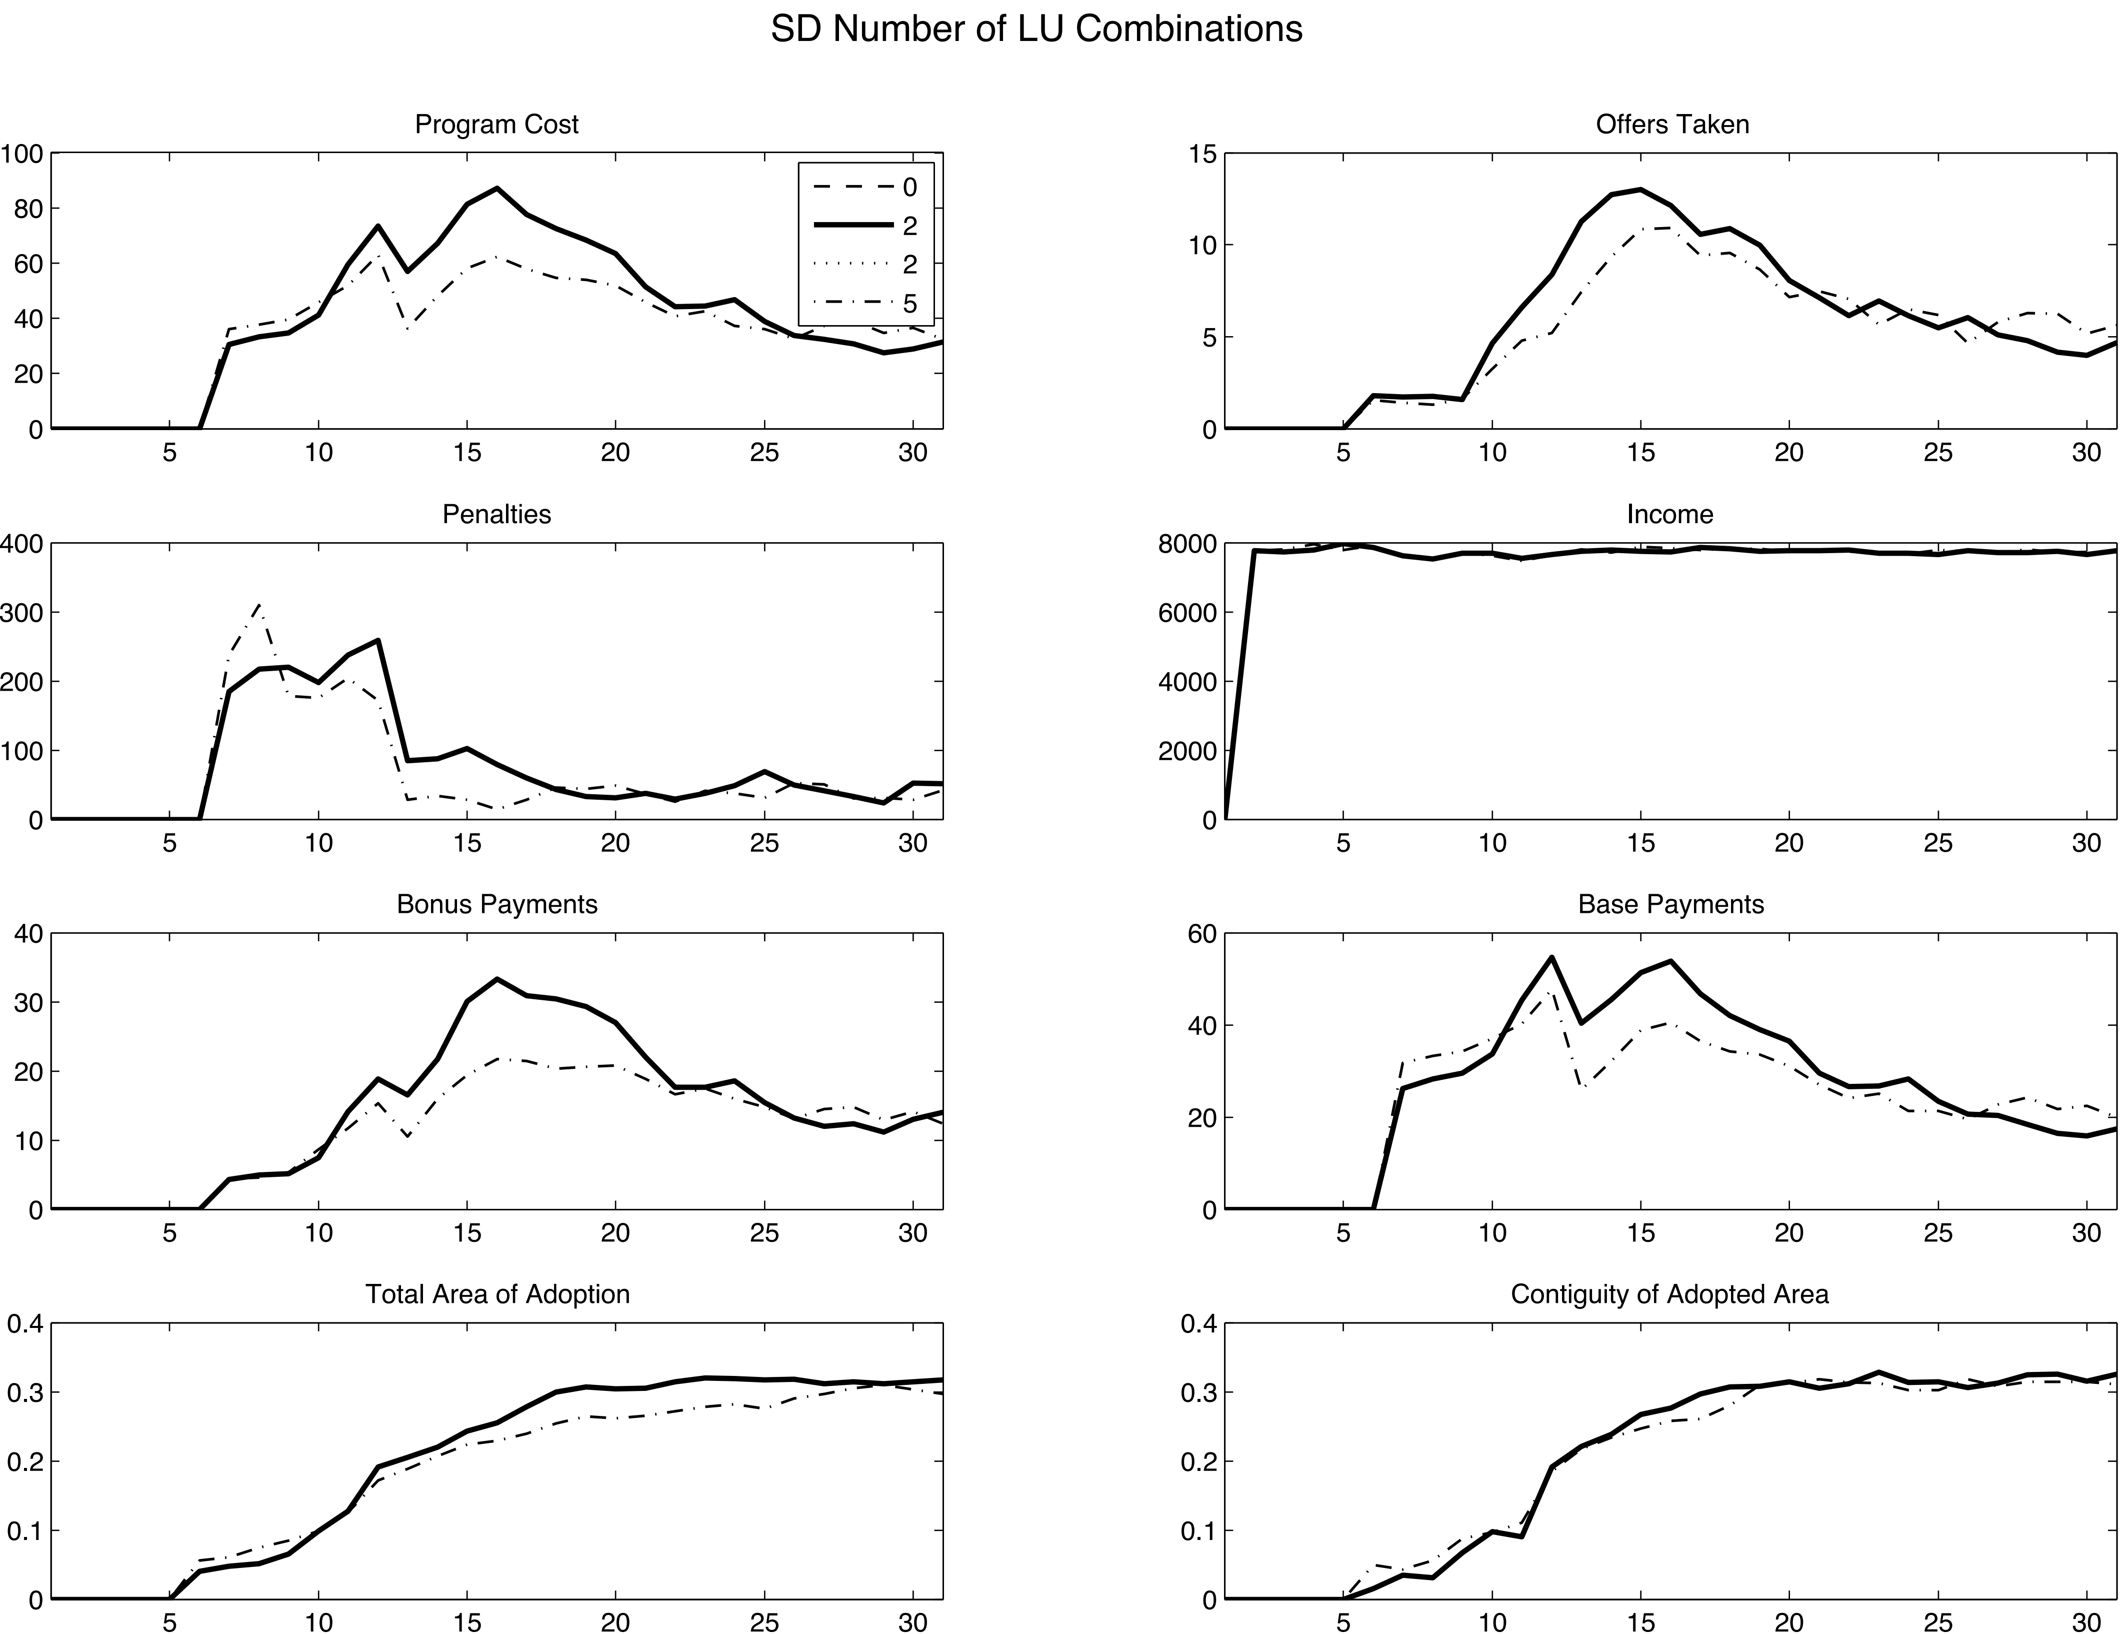
SD Number of LU Combinations
Program Cost
0
2
2
5
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

SD Plots per Farm

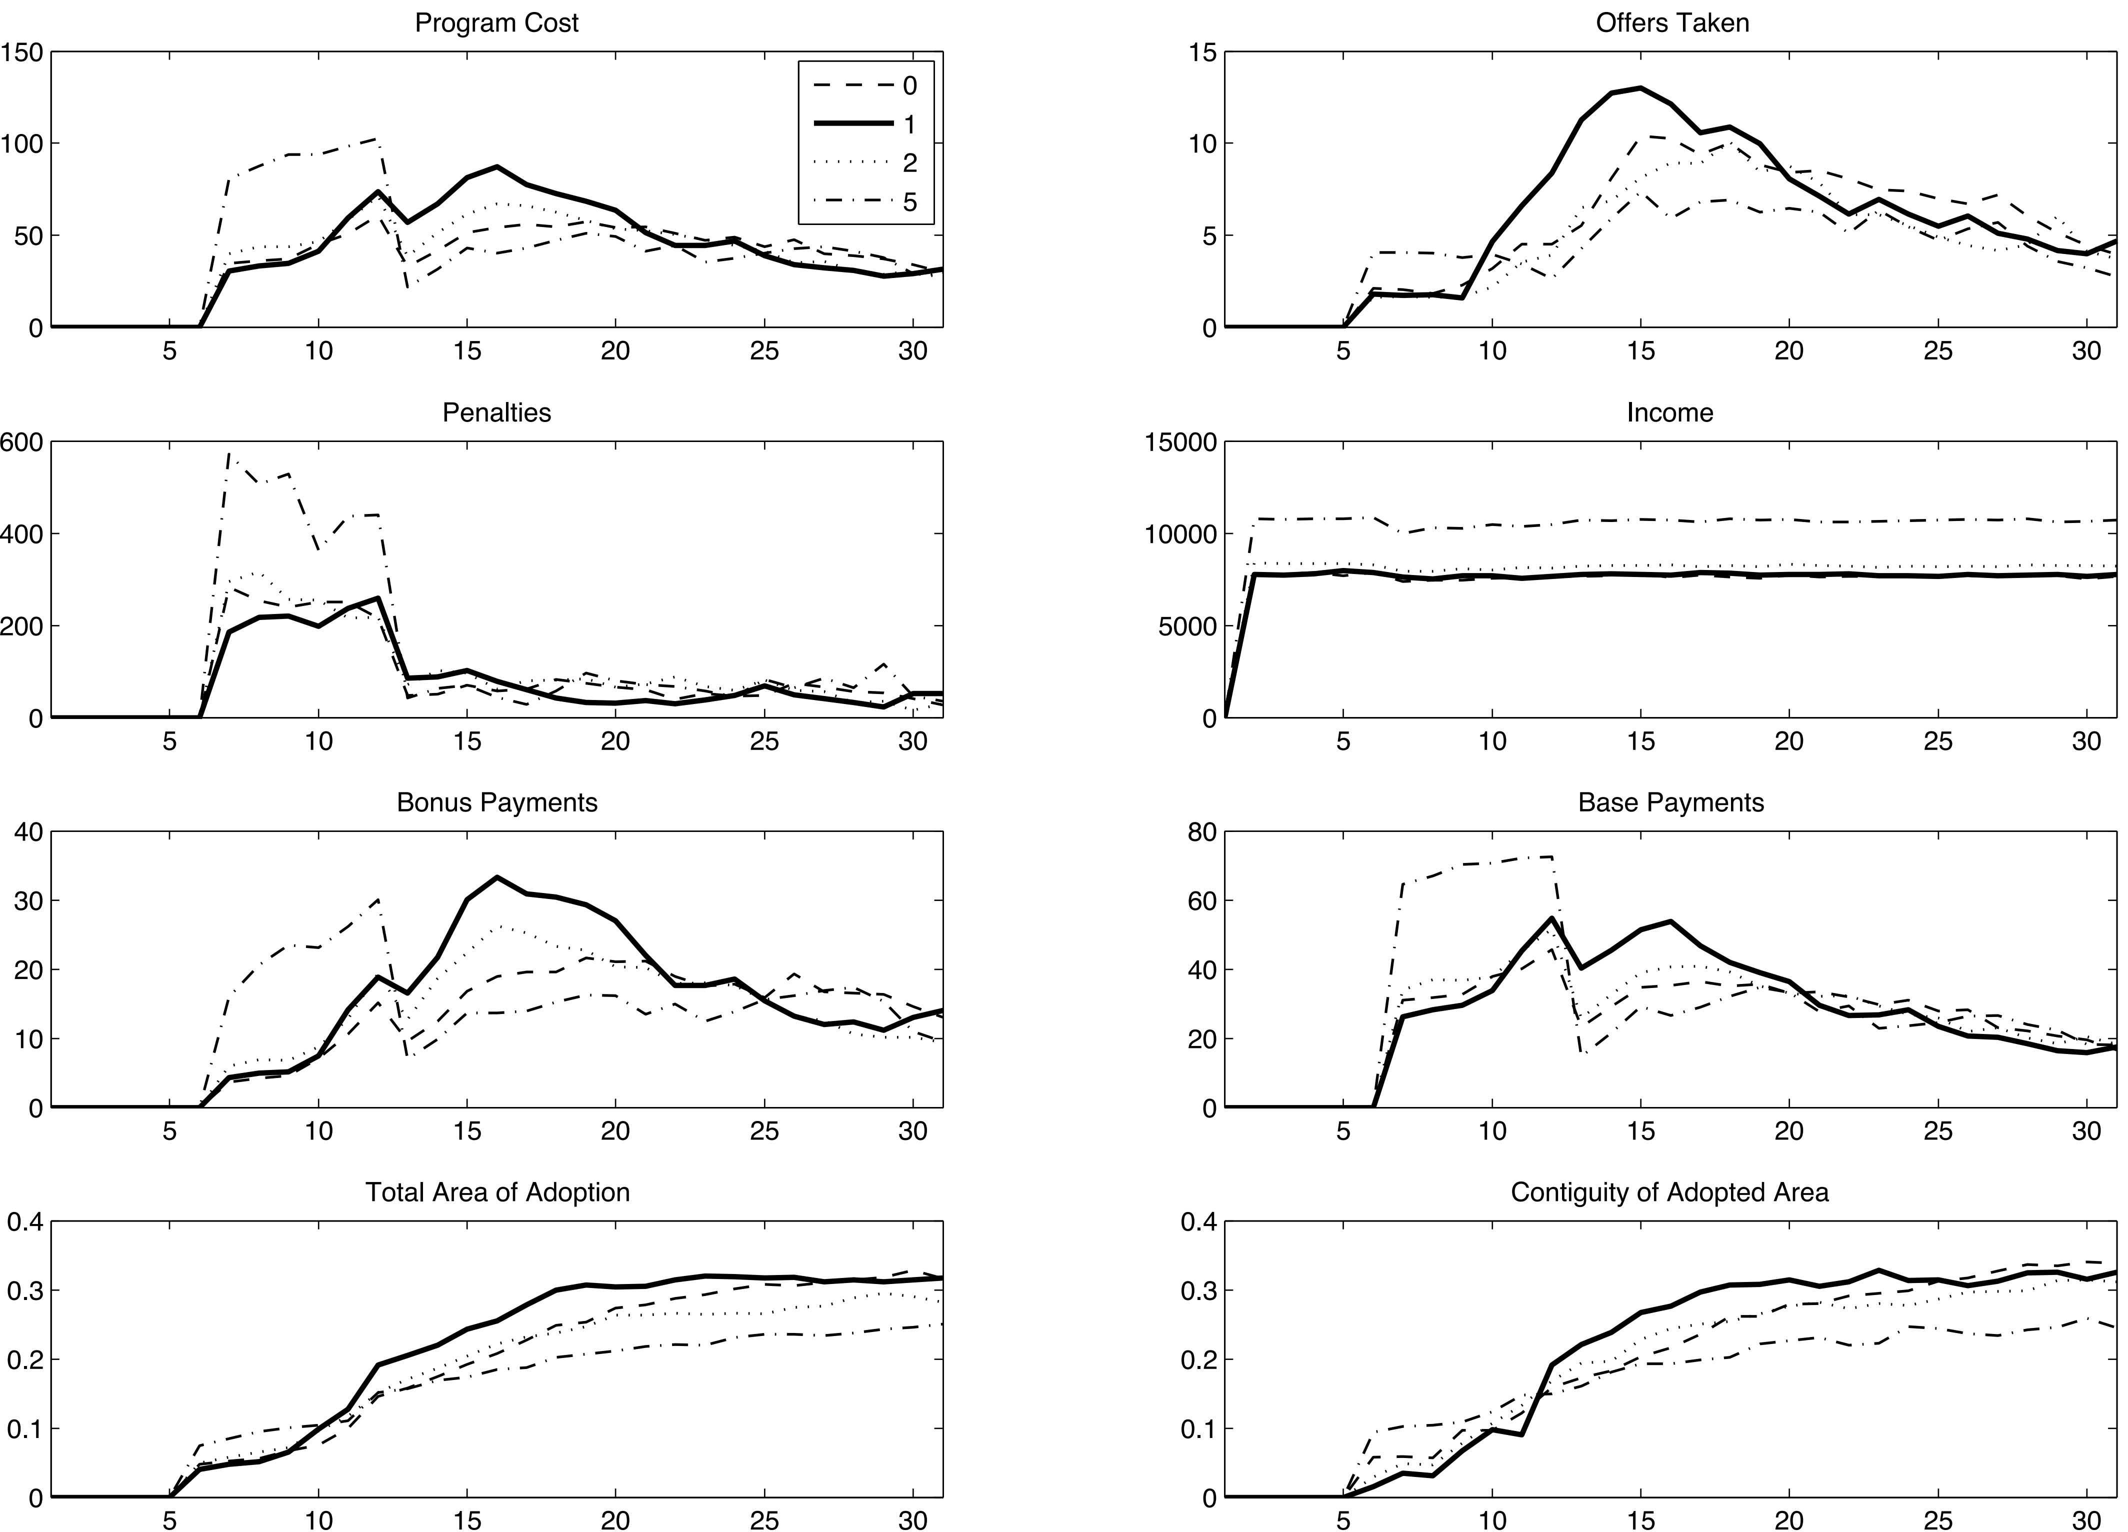

SD property size

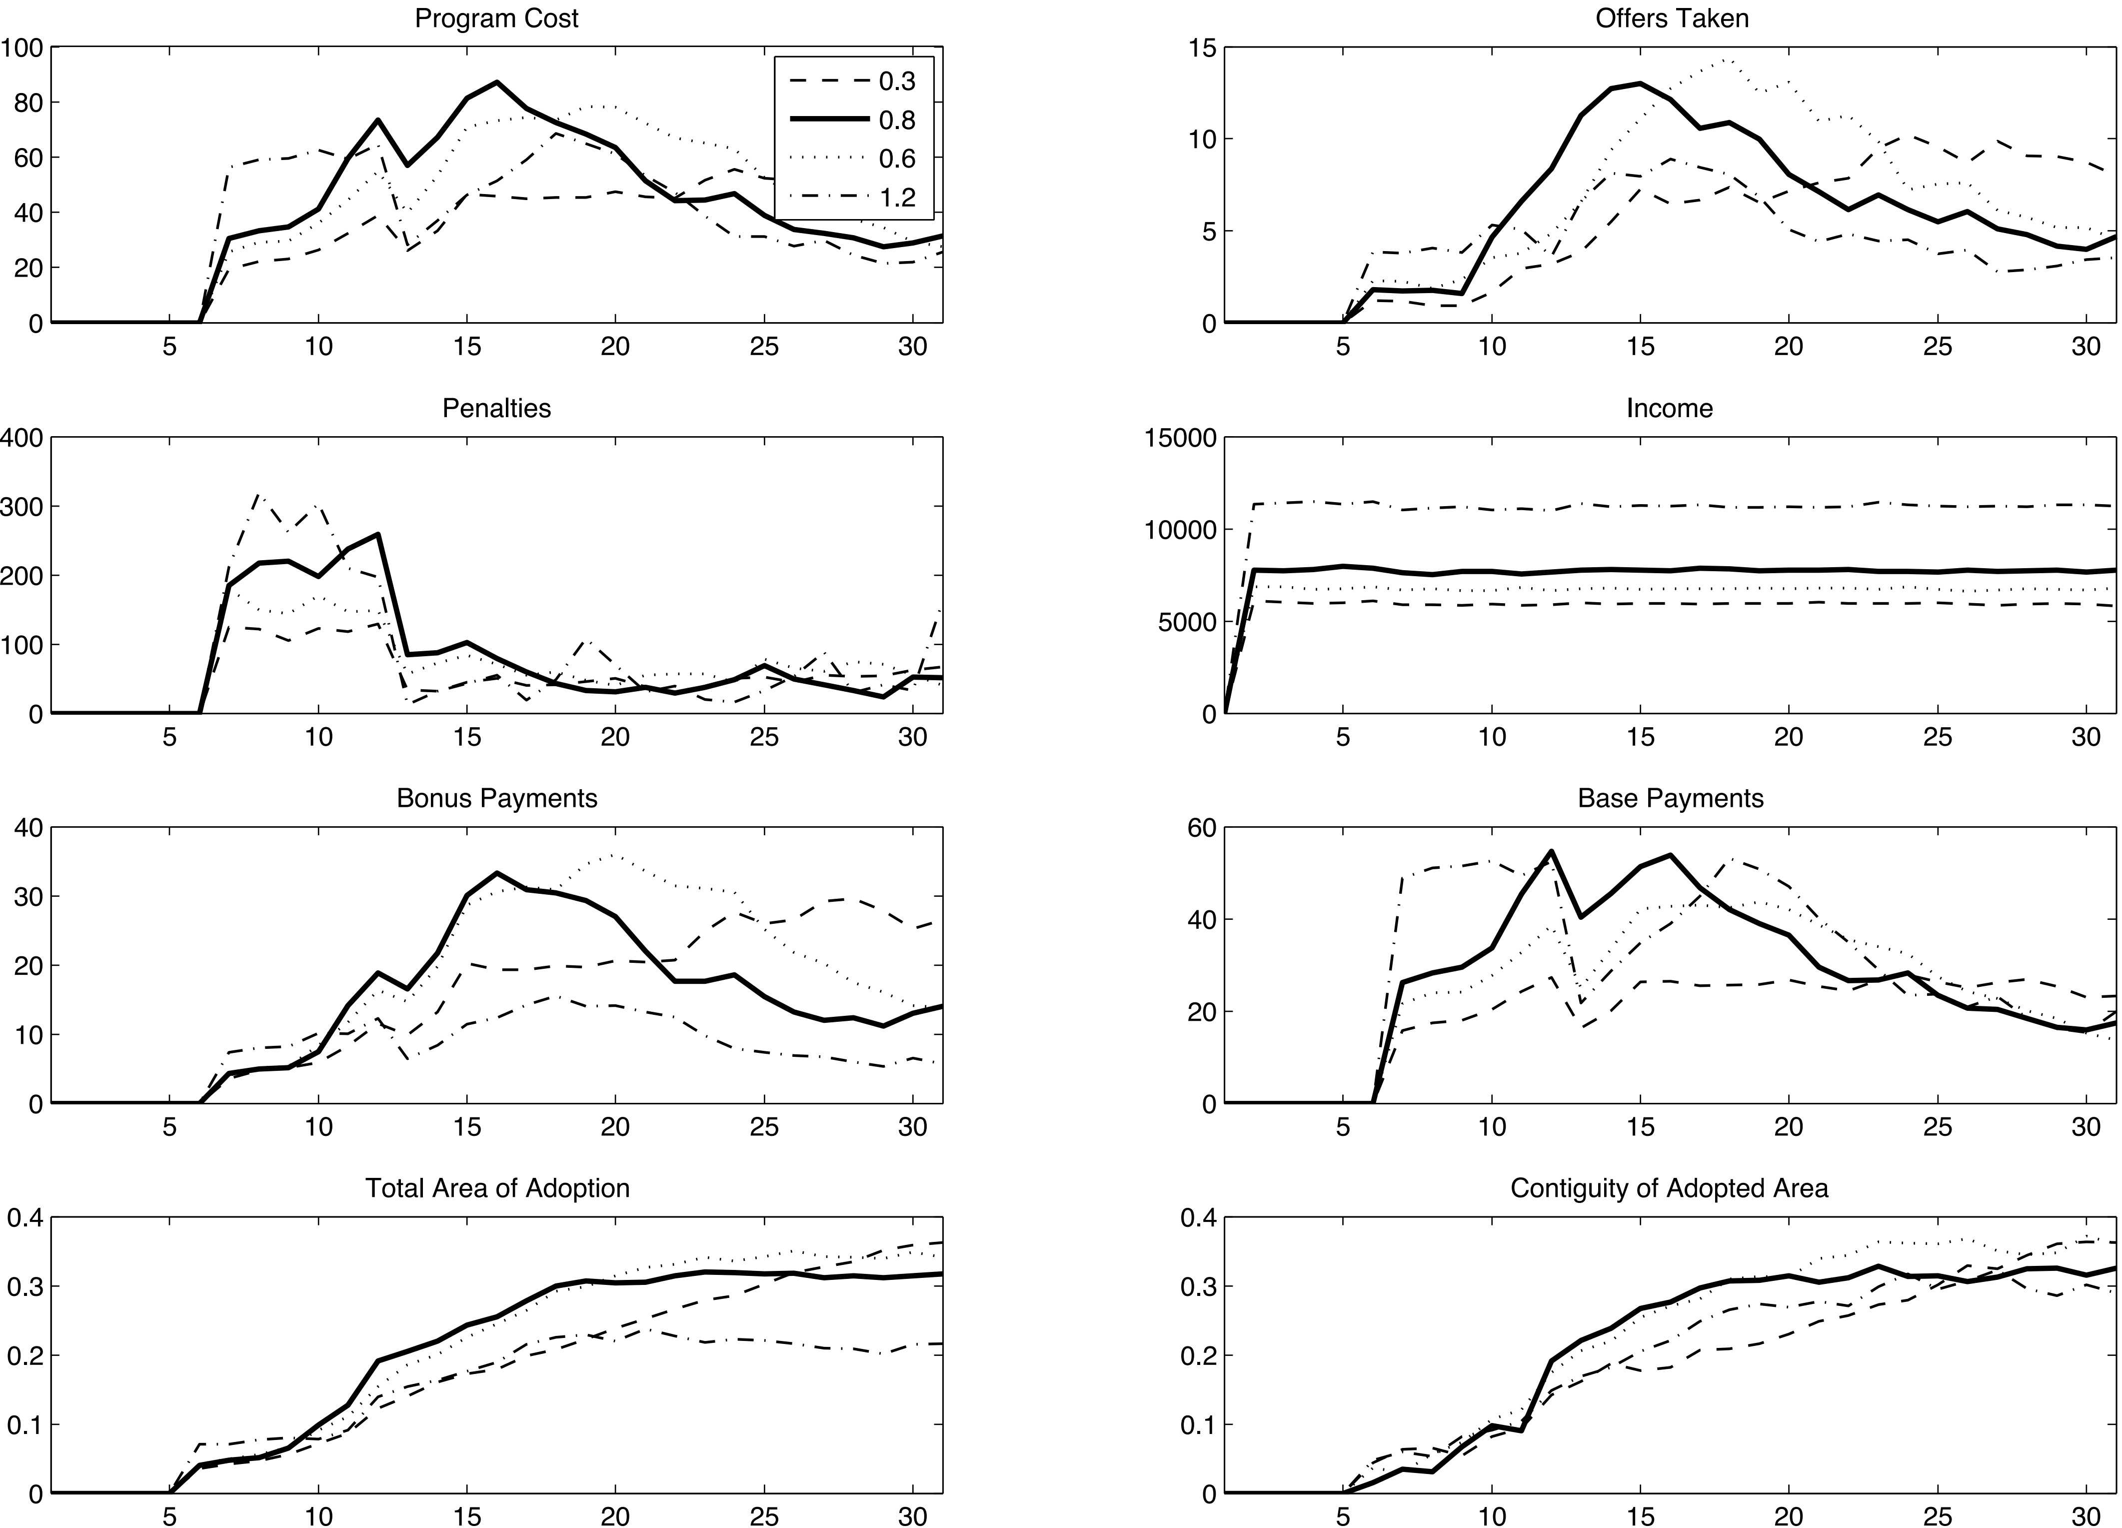

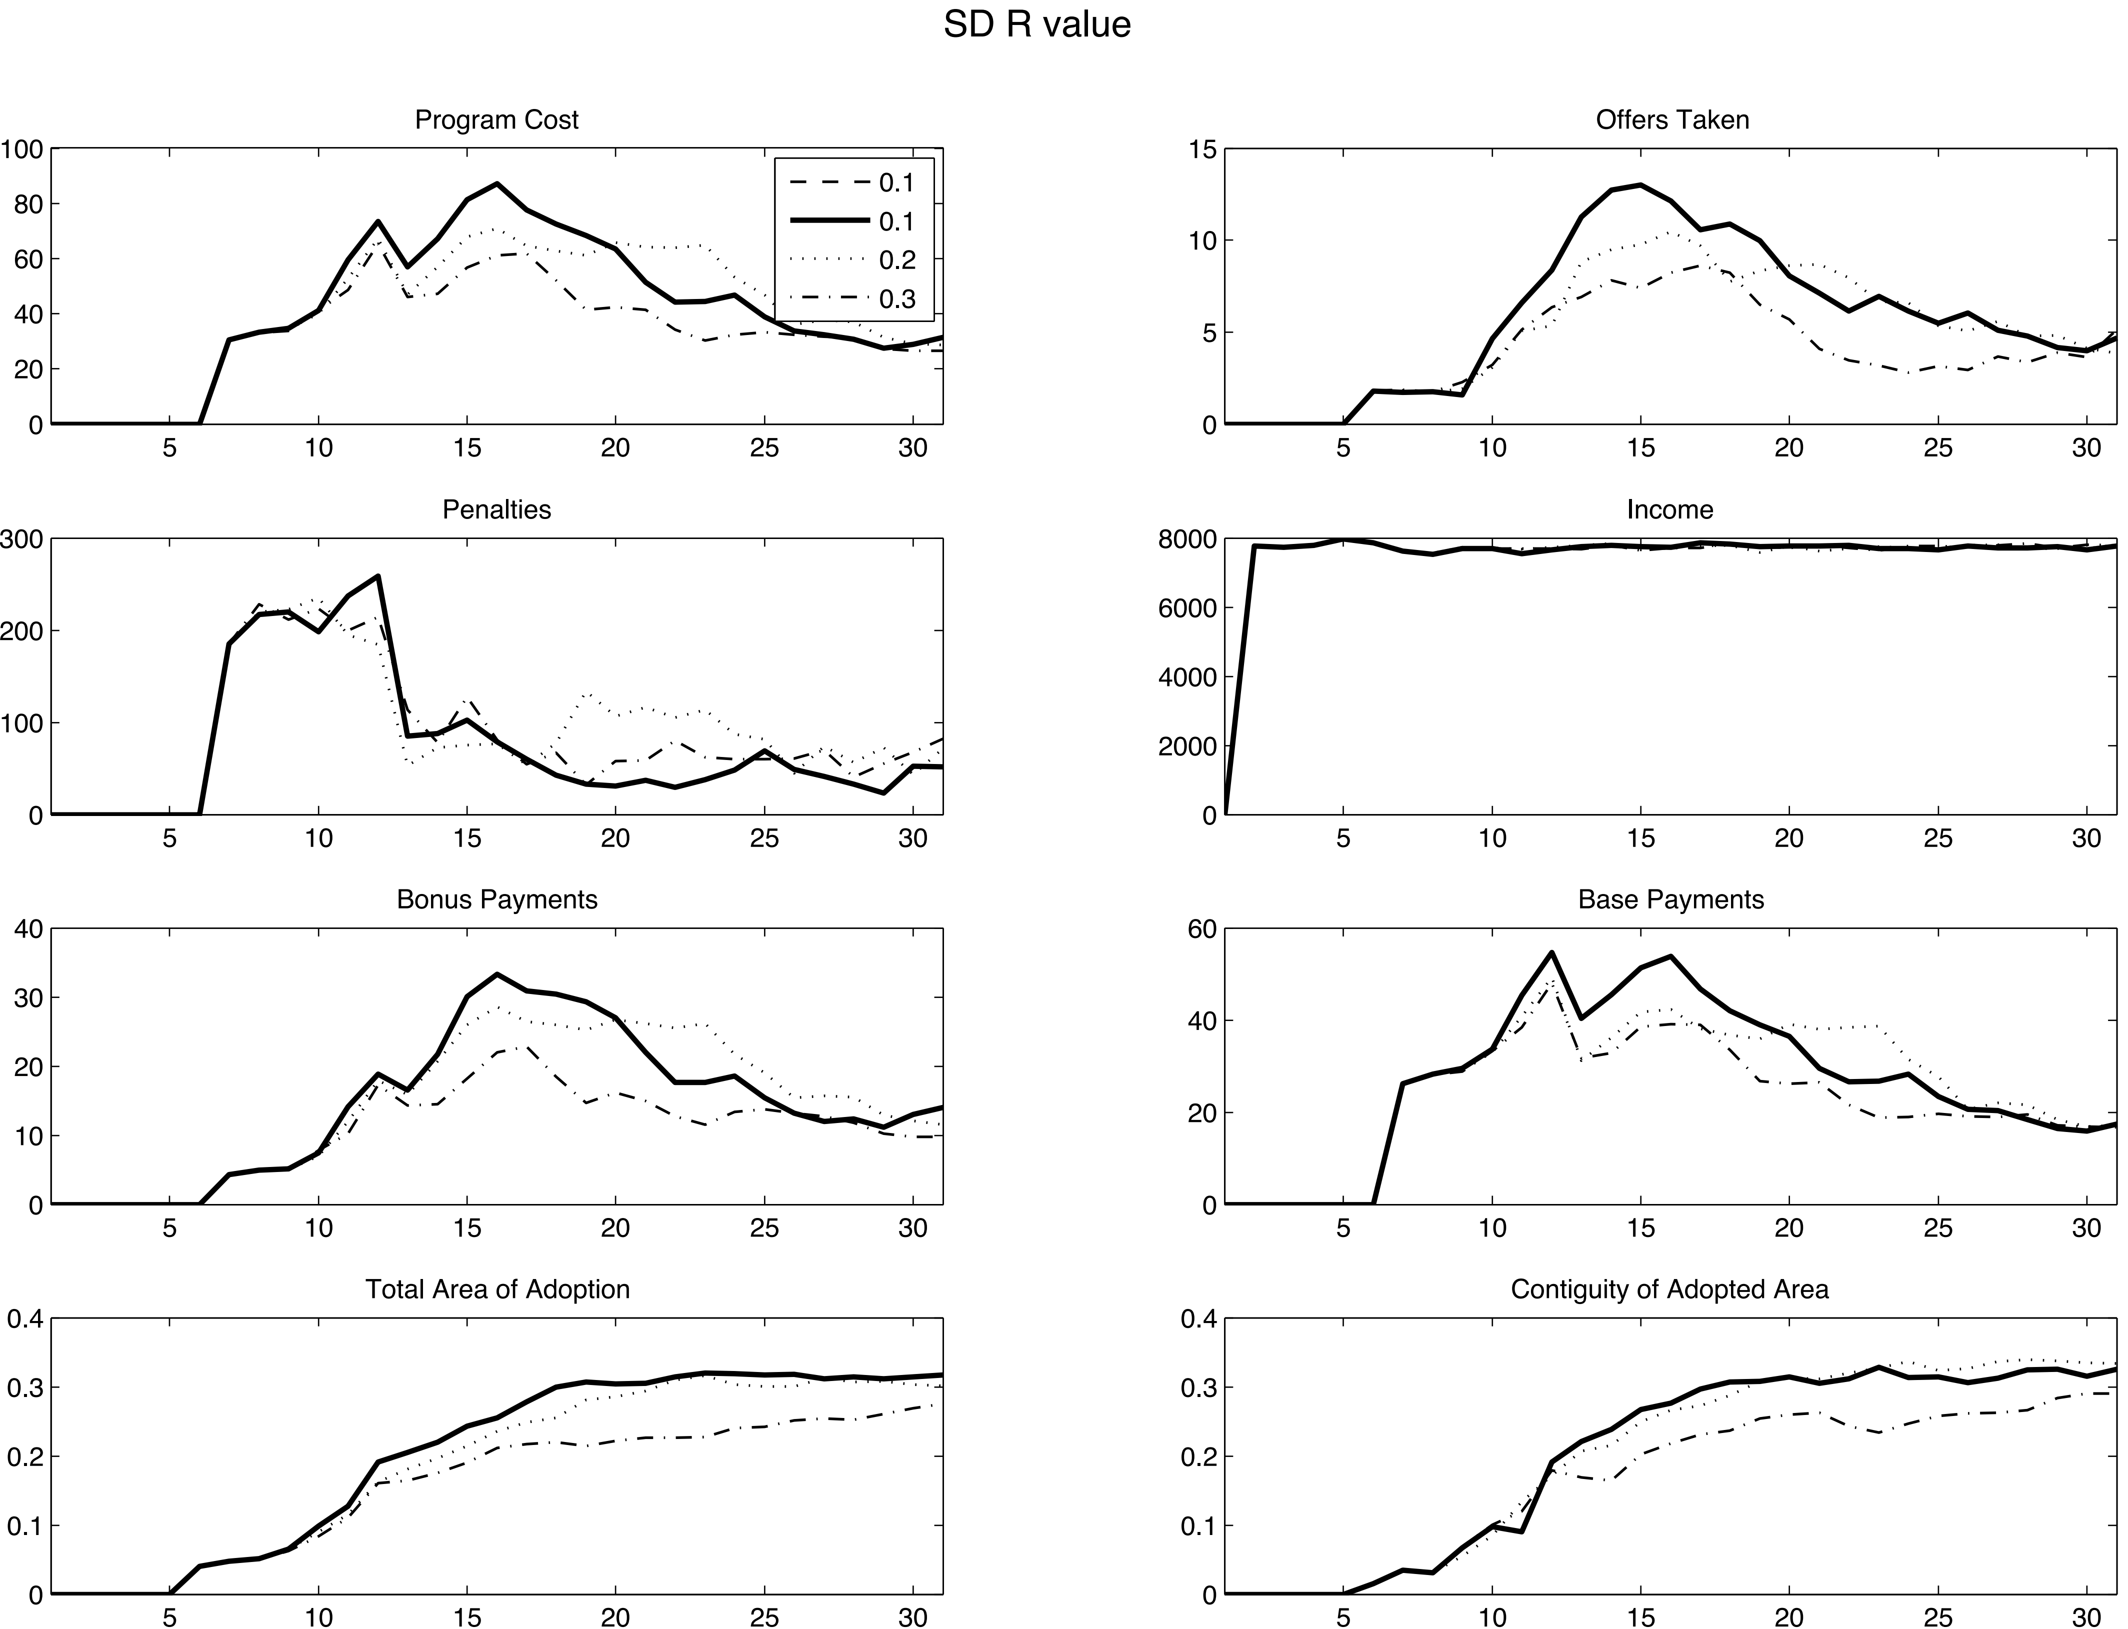
SD R value
Program Cost
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area
0.1
0.1
0.2
0.3

SD soil quality–farmer efficiency

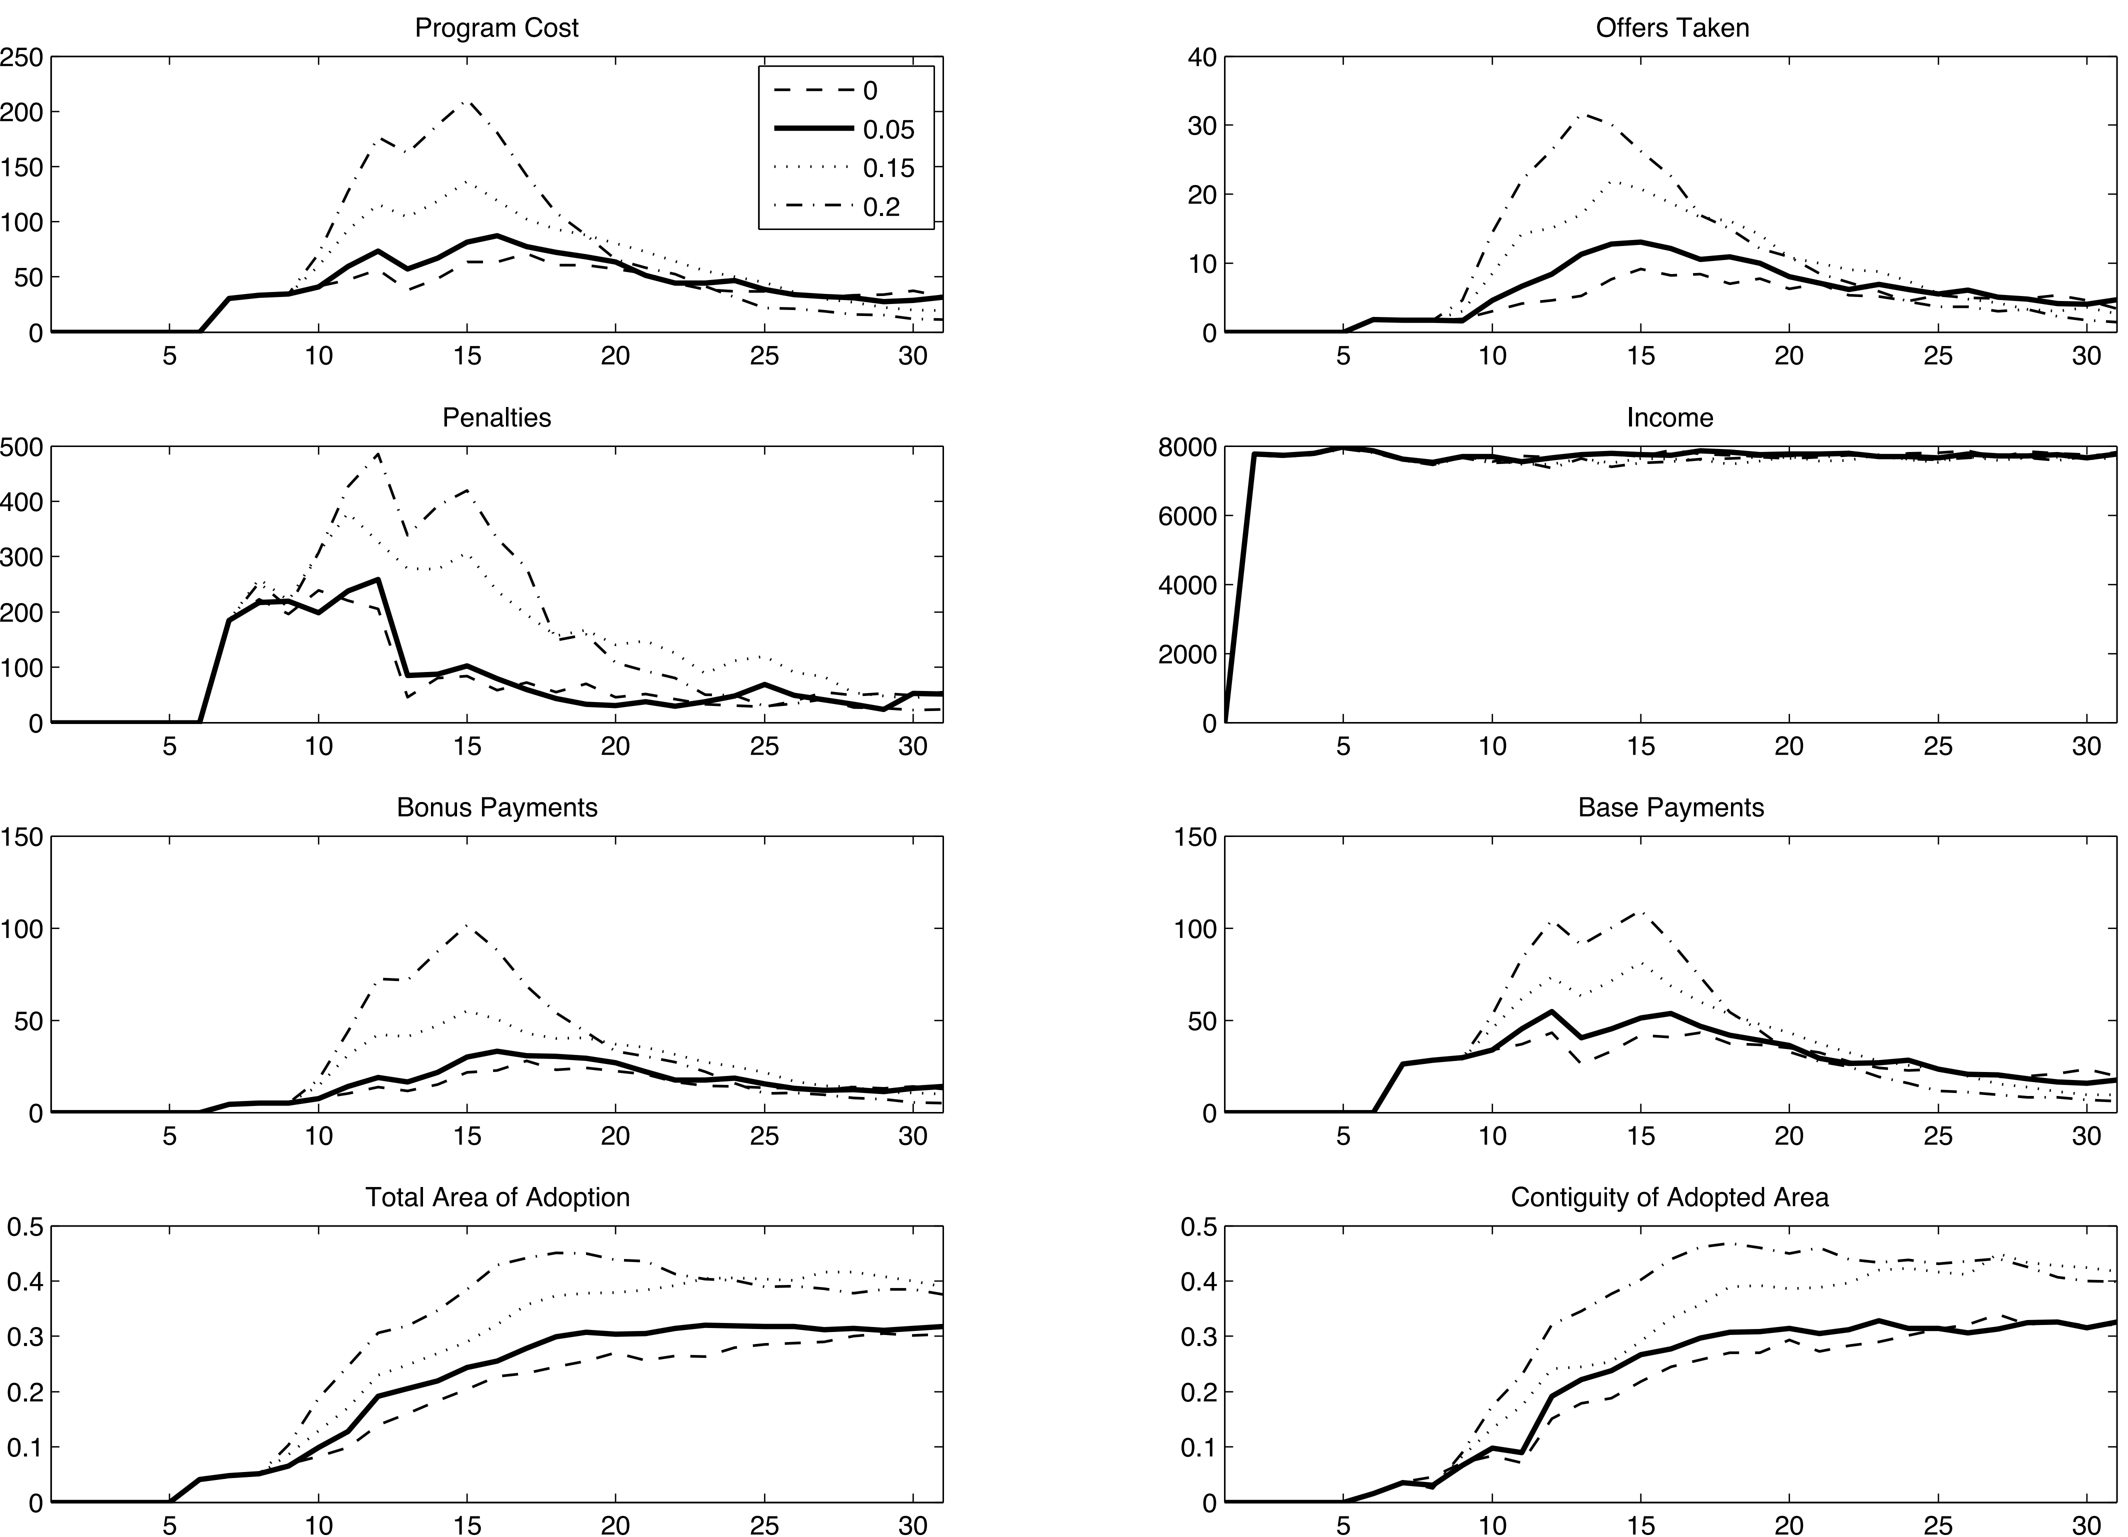

Start of early adoption period

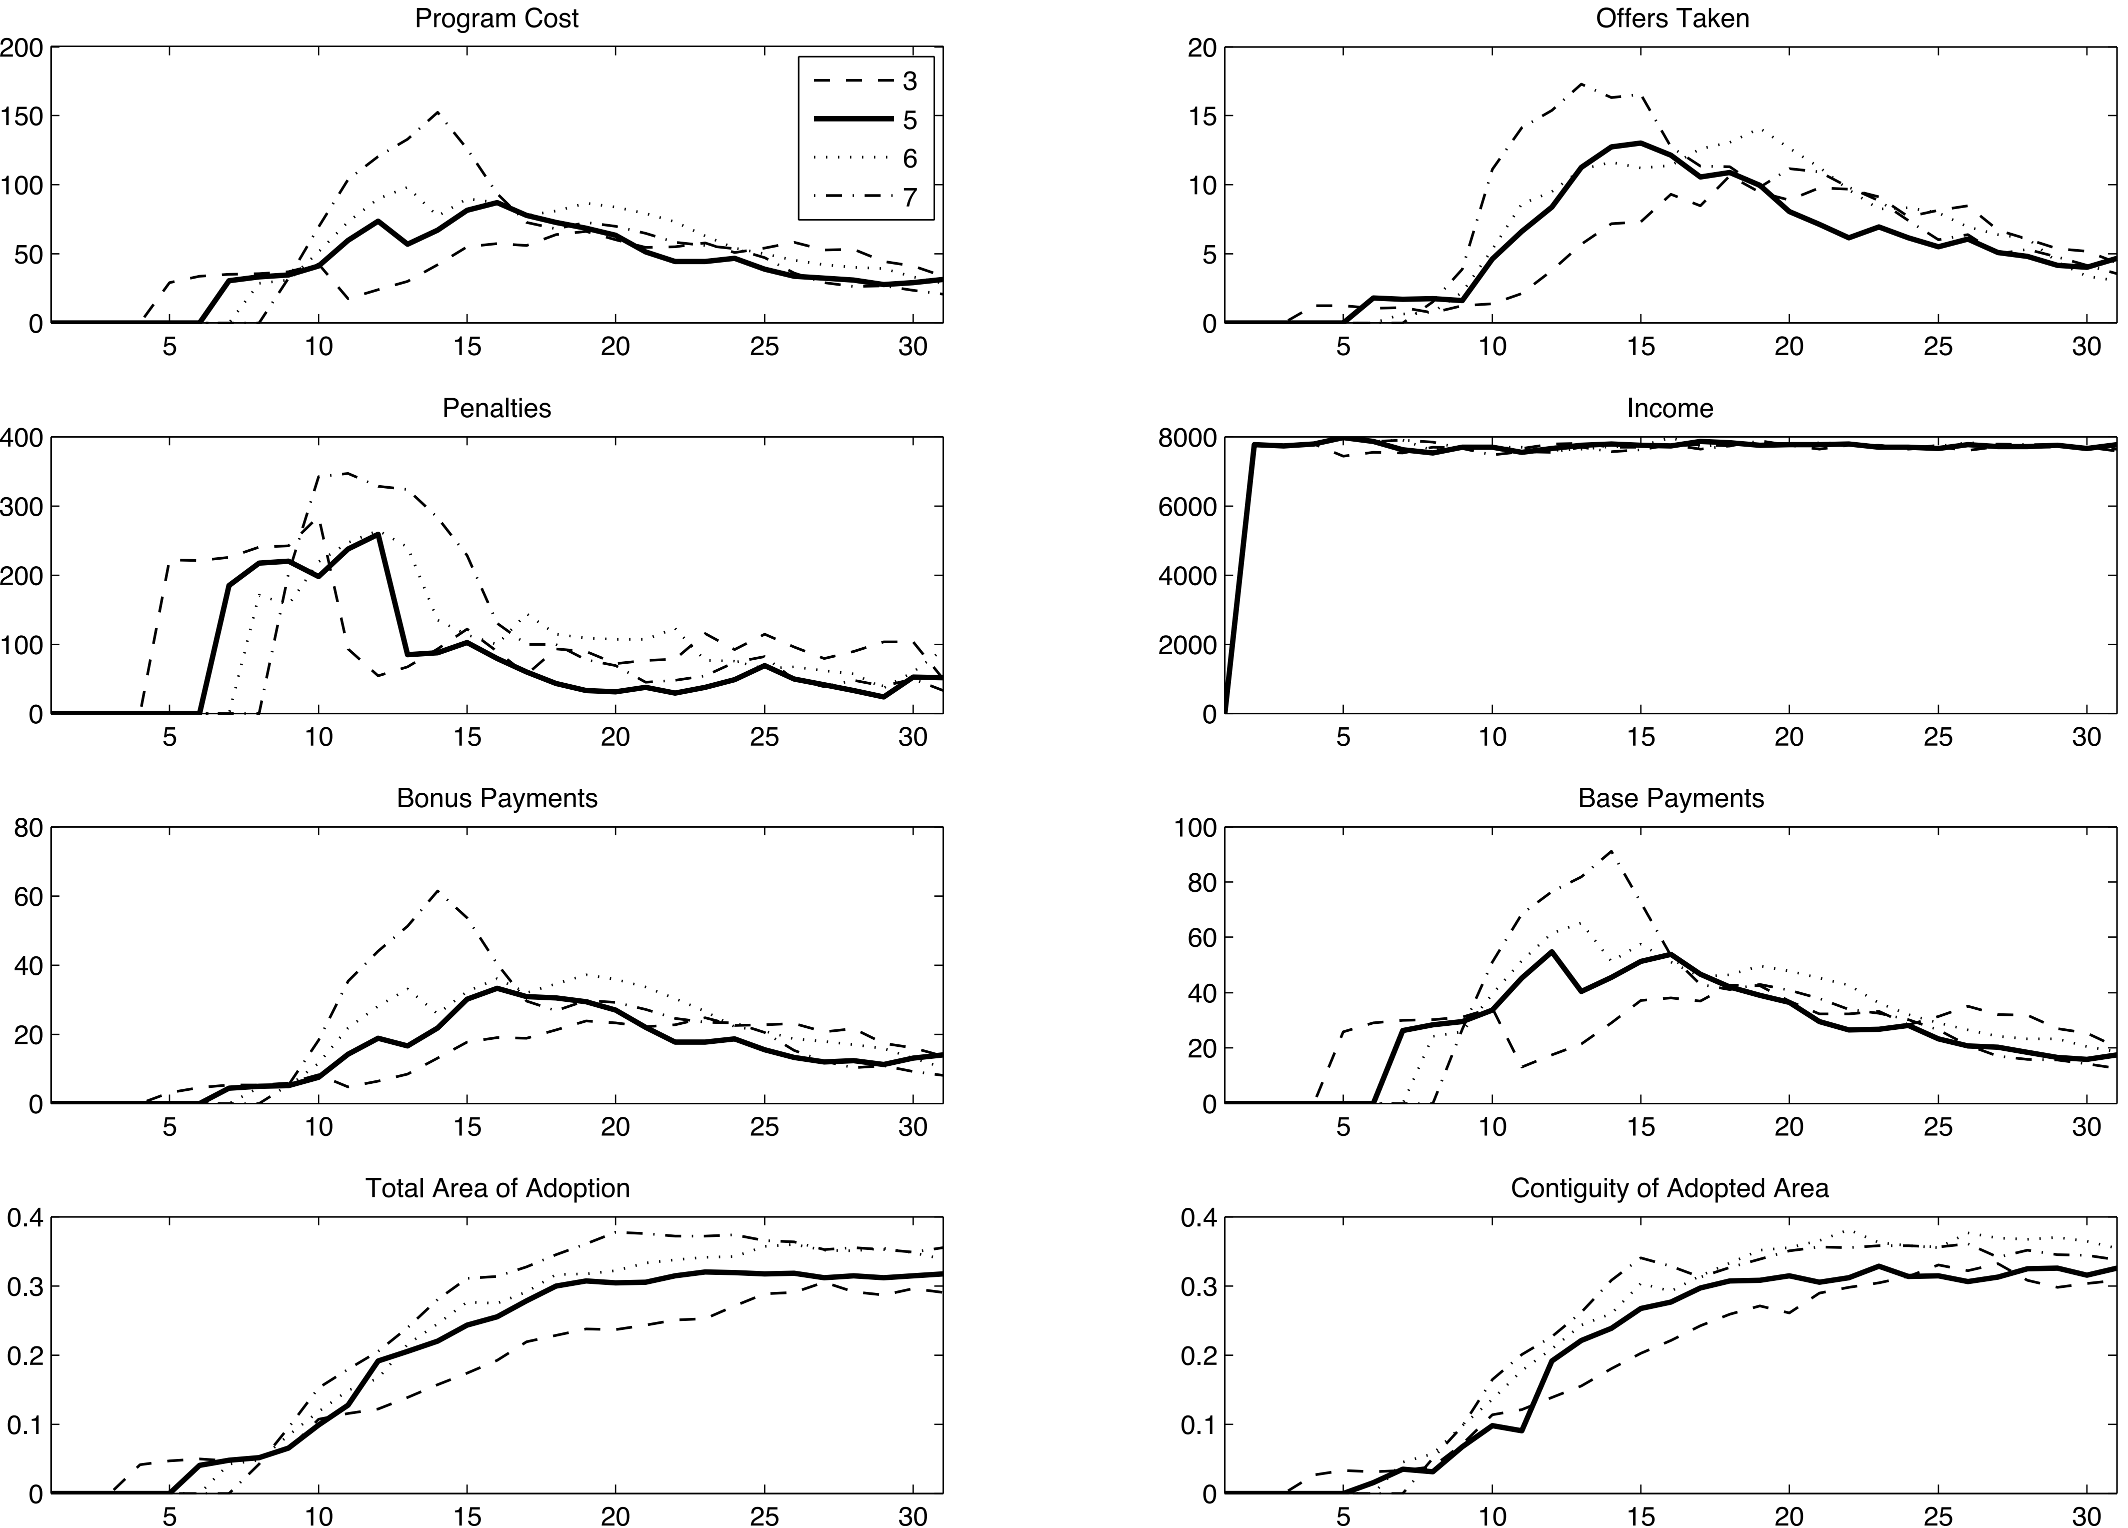

Transaction cost to Conservation Ag

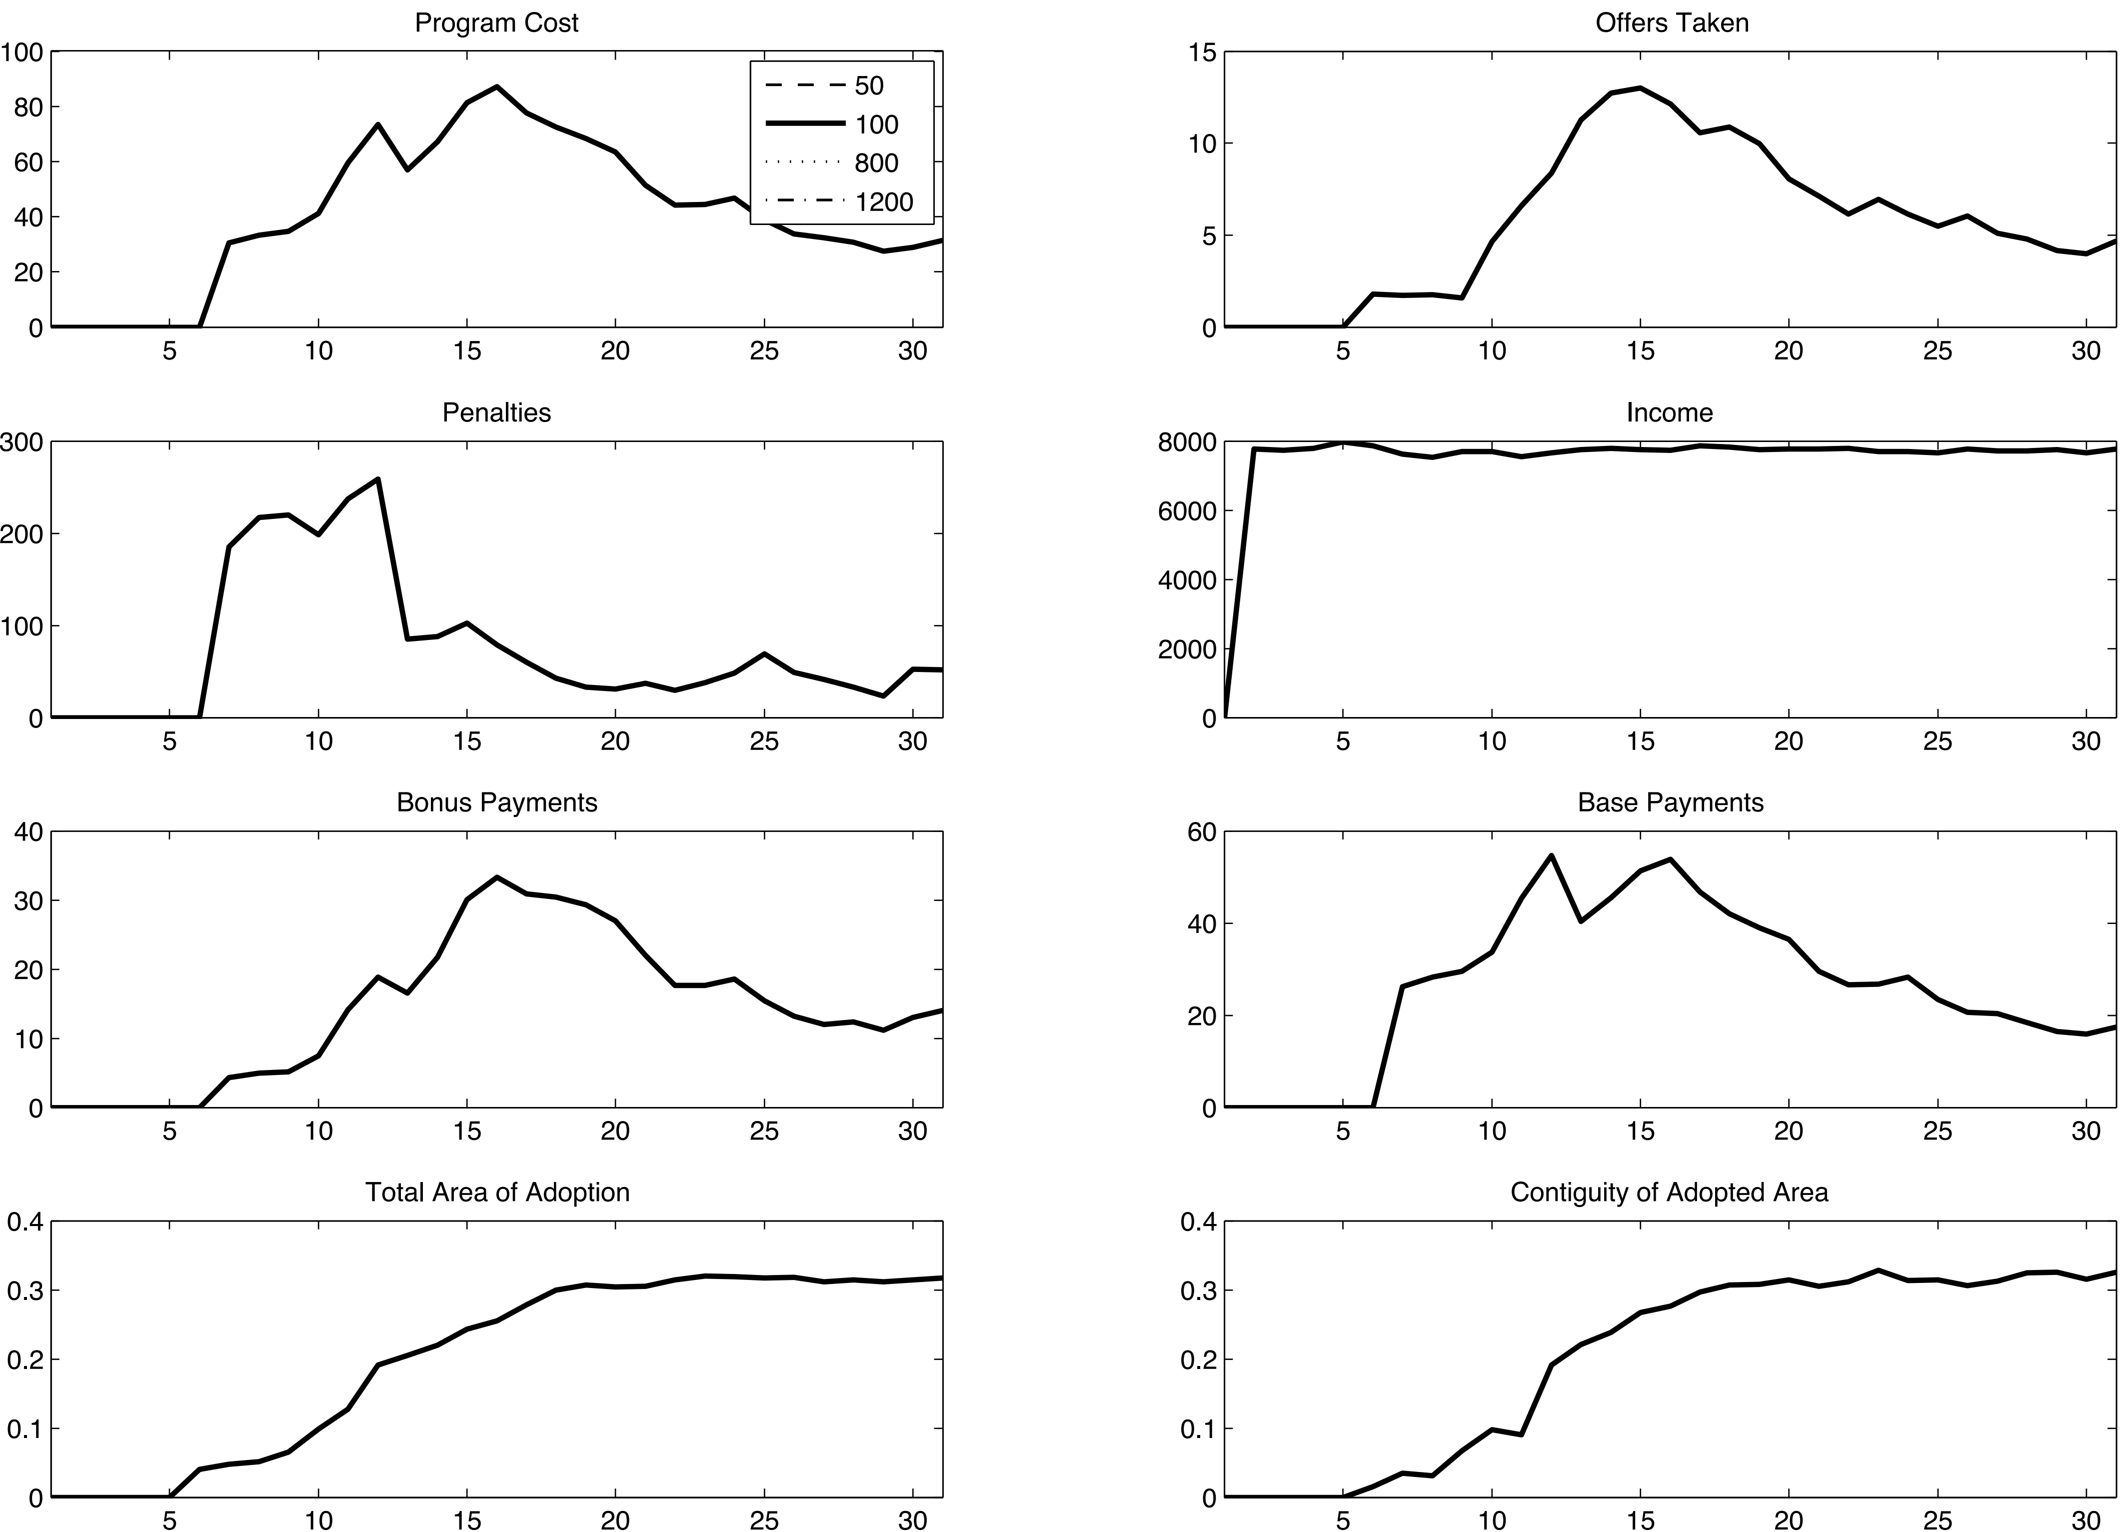

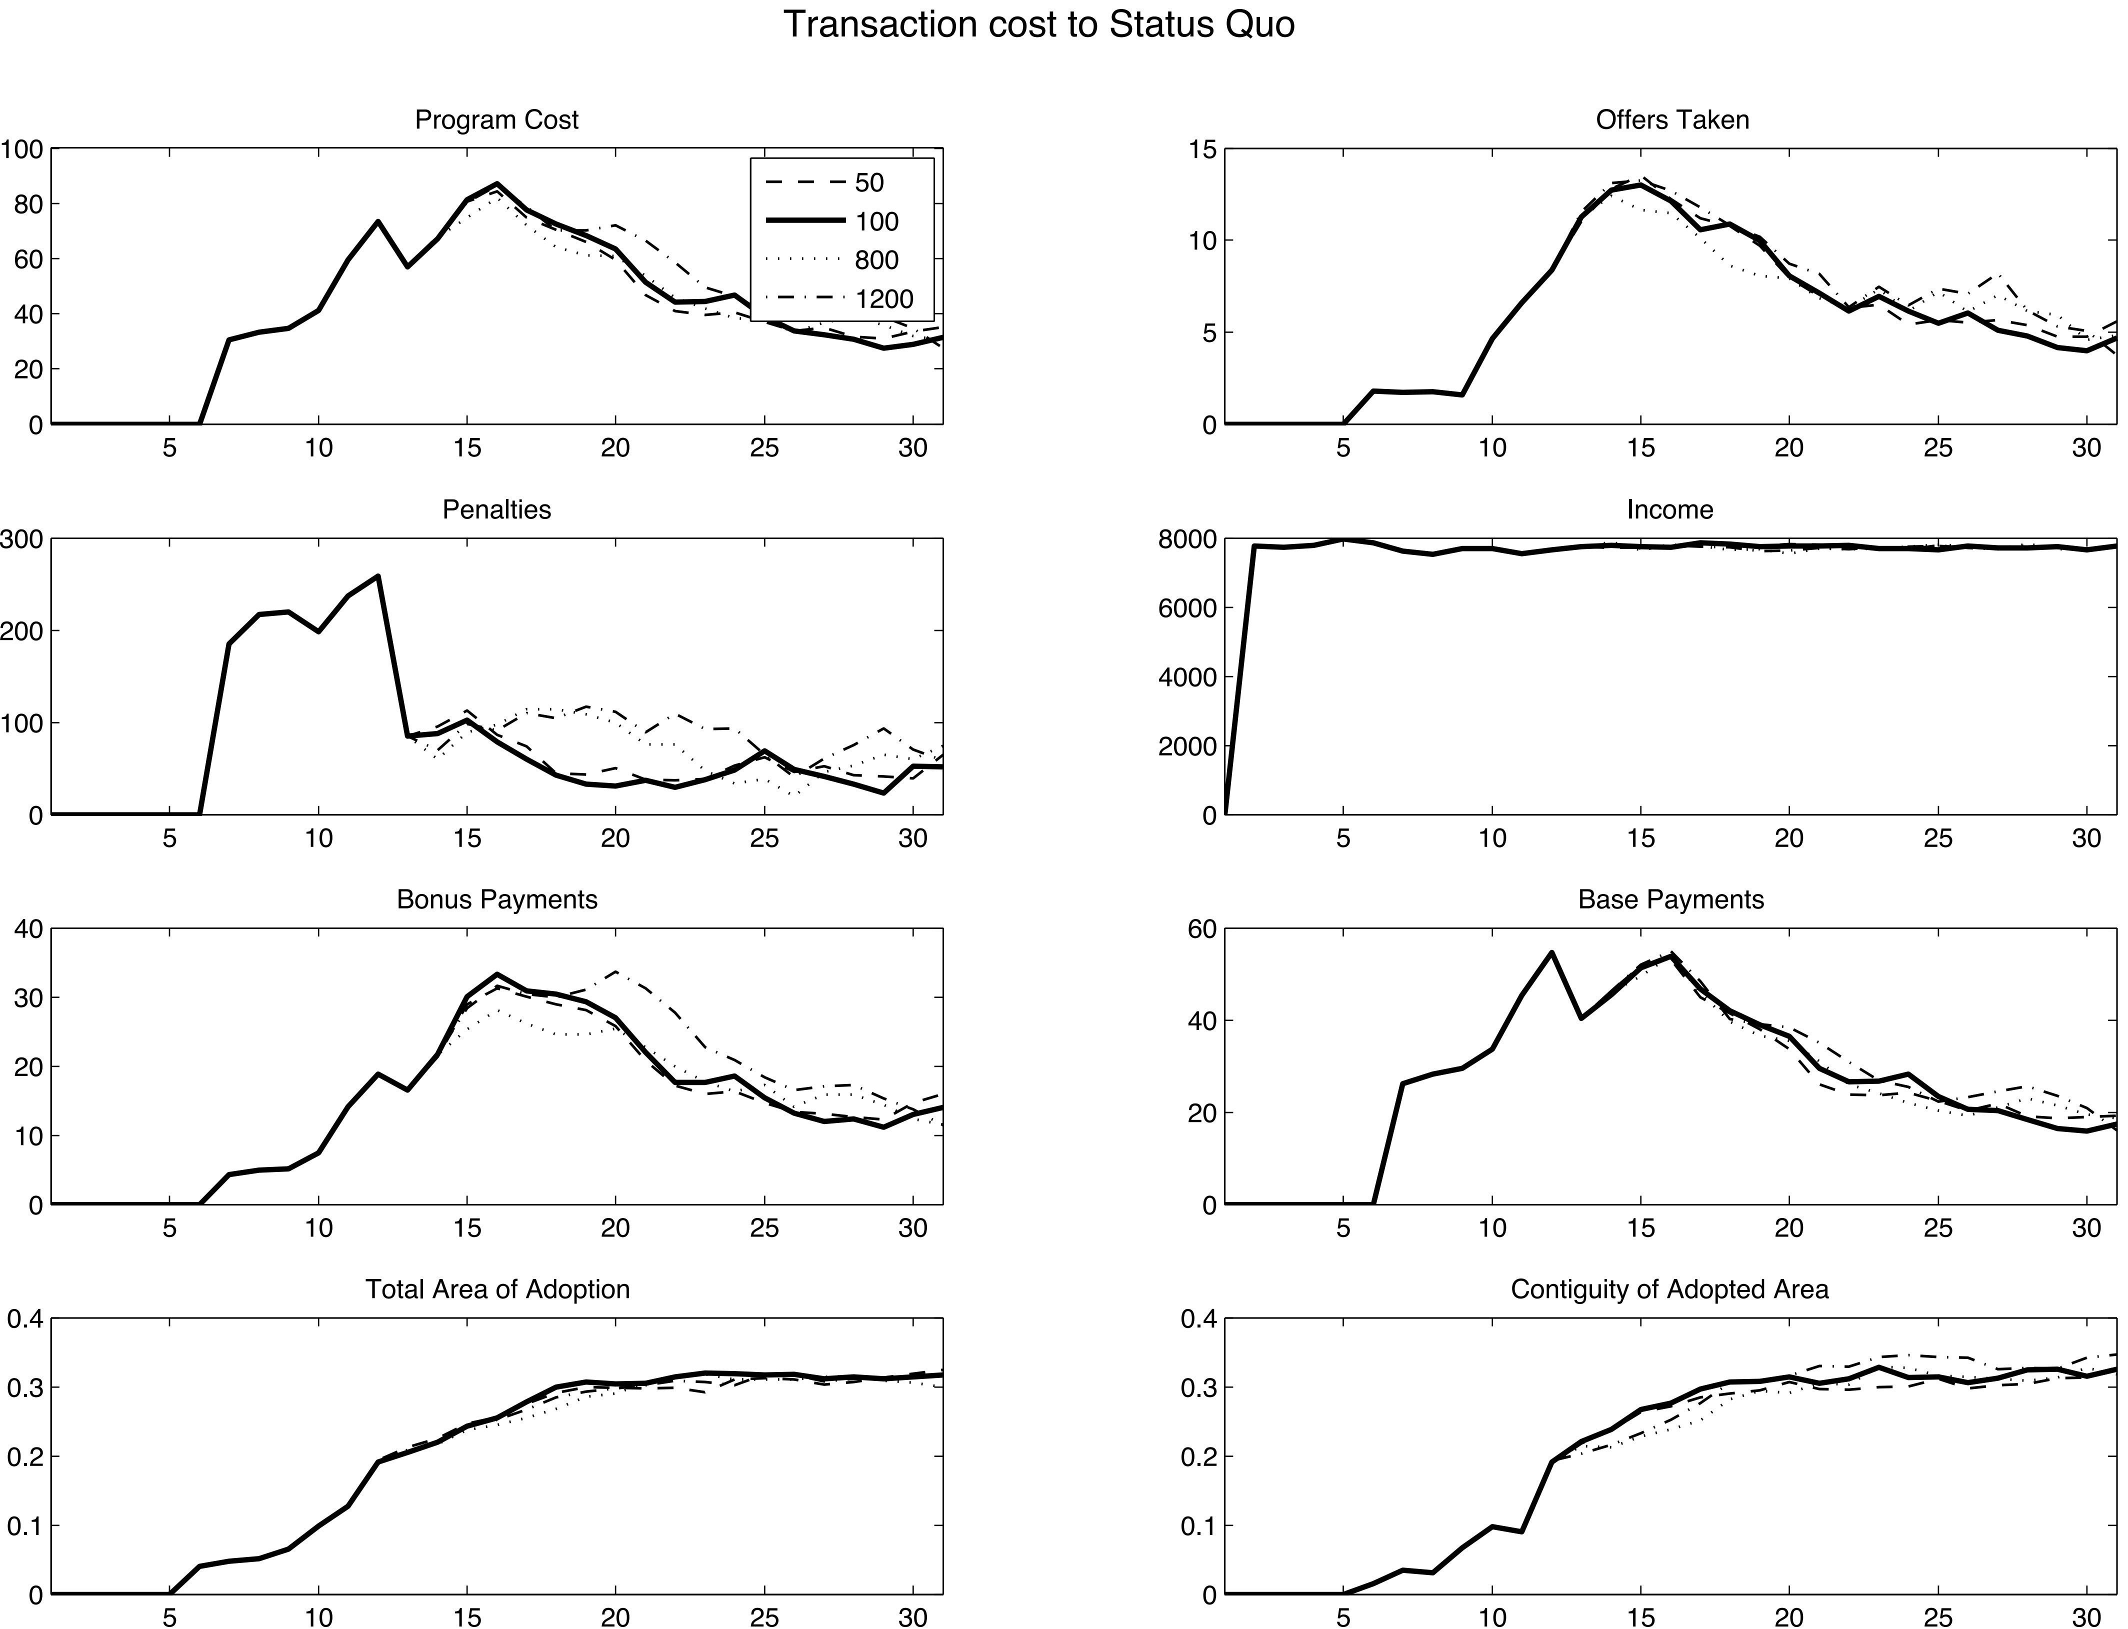
Transaction cost to Status Quo
Program Cost
50
100
800
1200
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

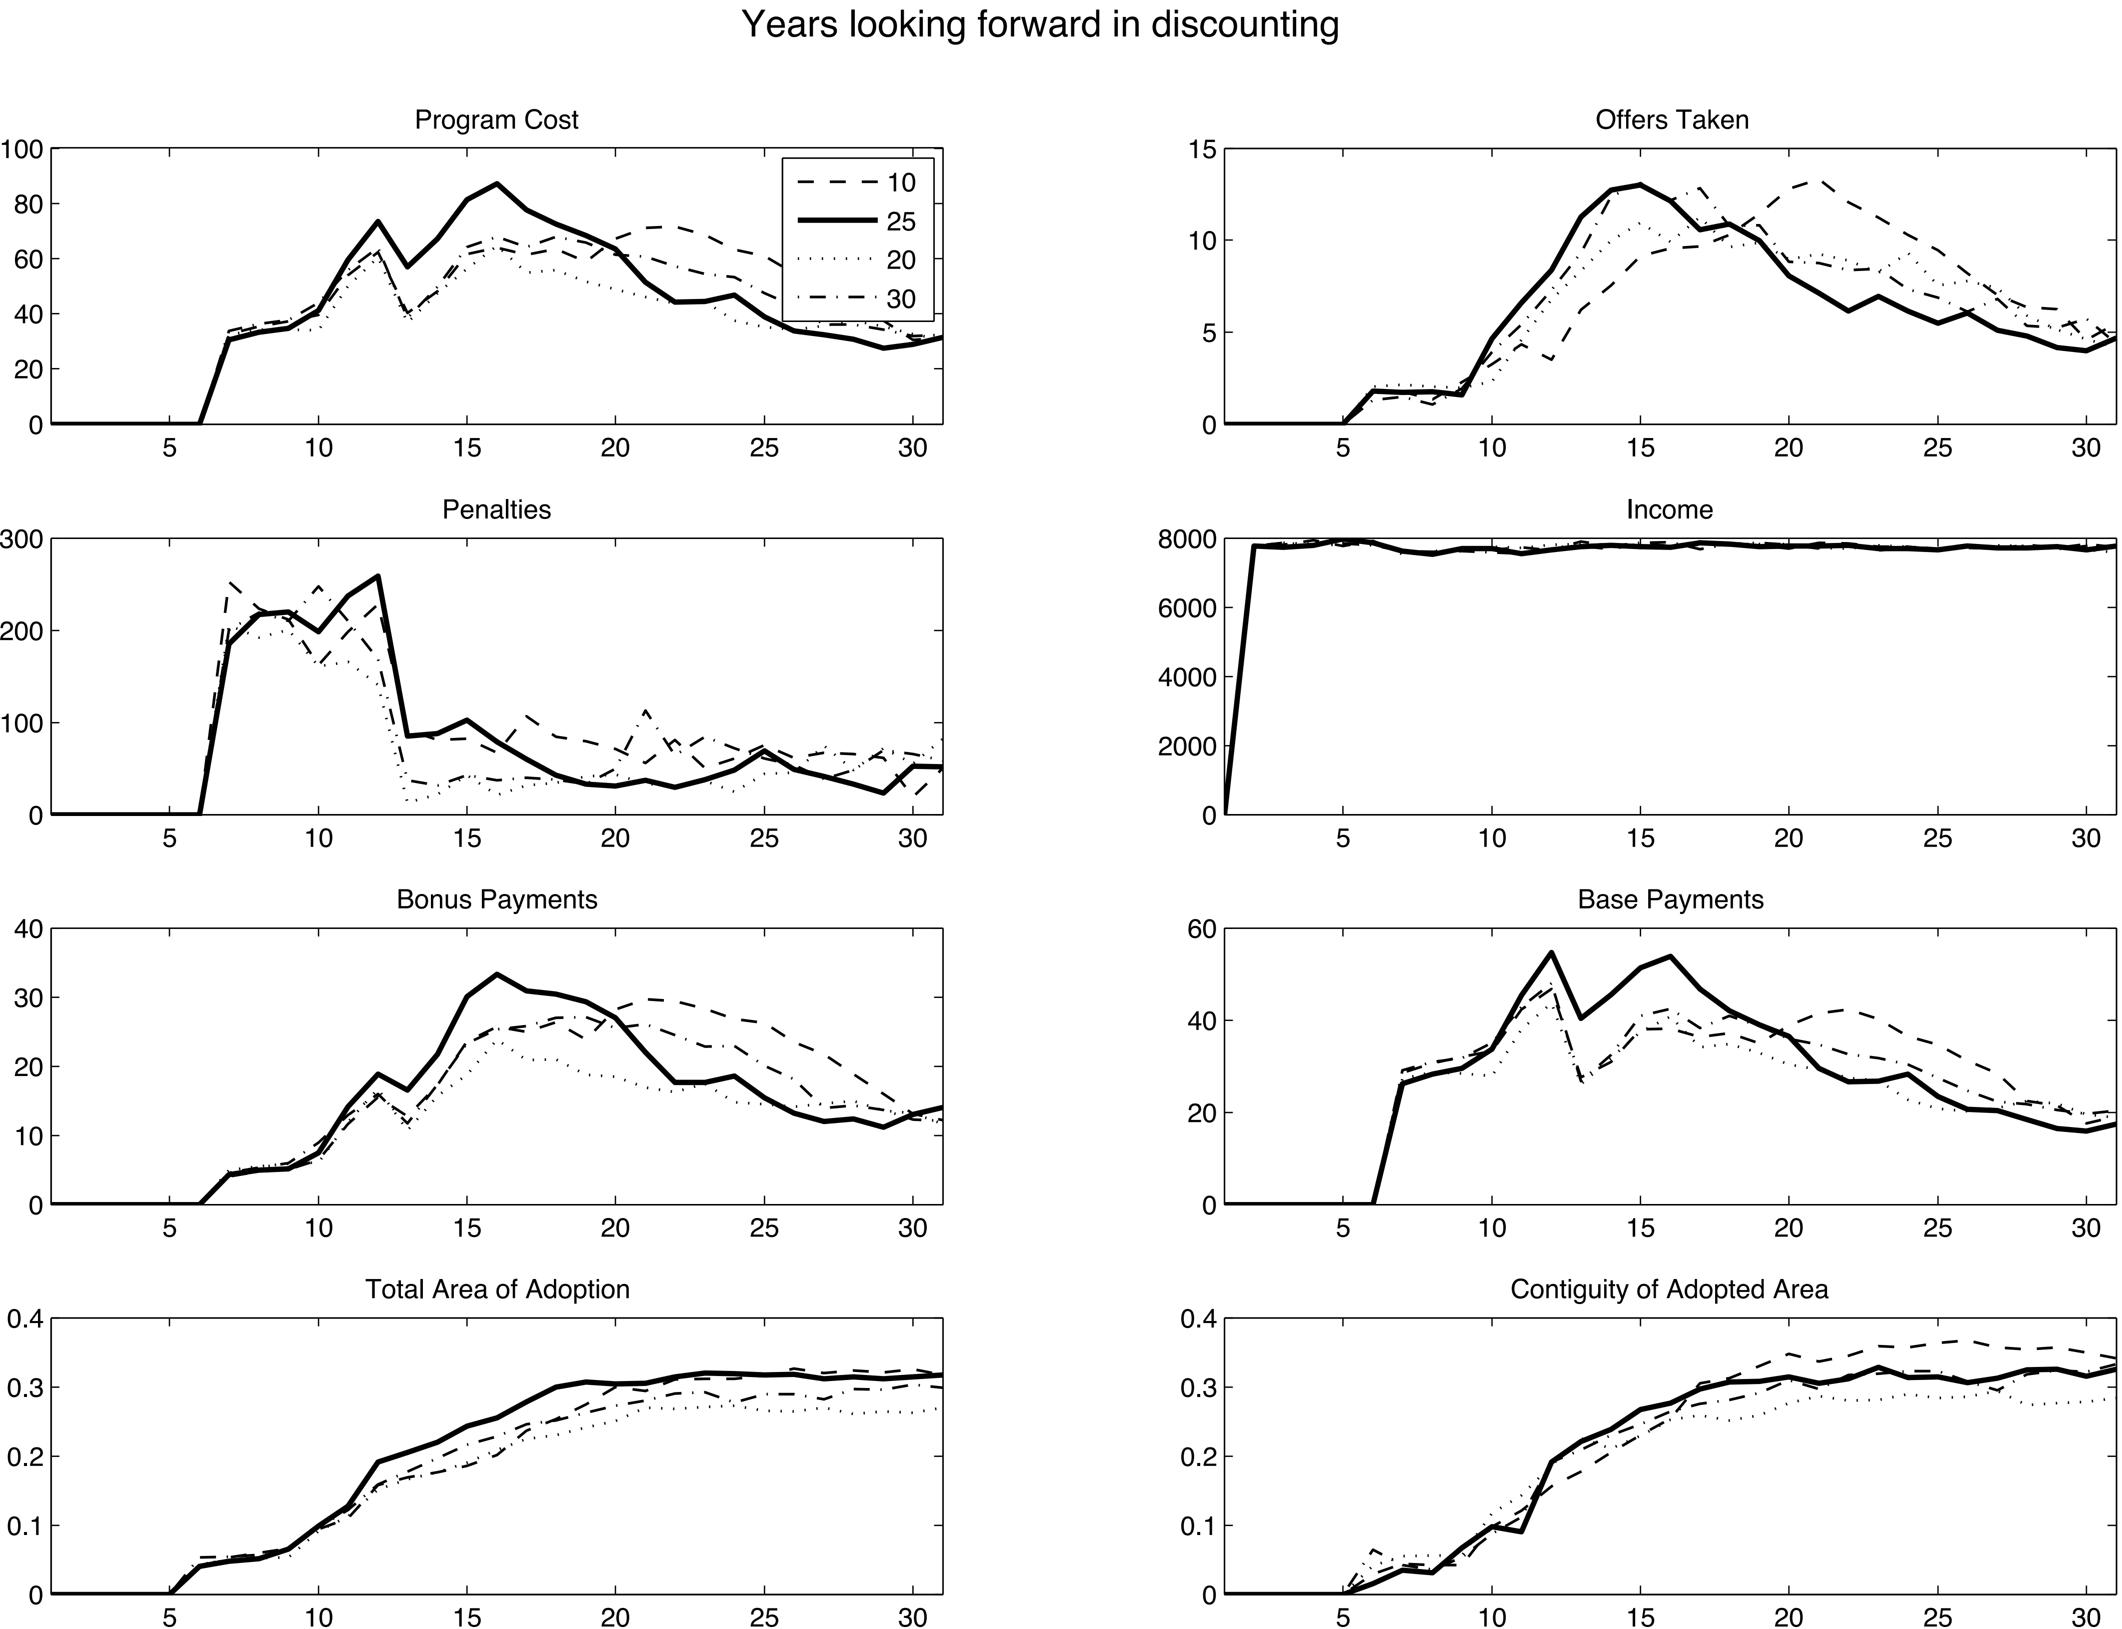

Years looking forward in discounting
Program Cost
10
25
20
30
Offers Taken
Penalties
Income
Bonus Payments
Base Payments
Total Area of Adoption
Contiguity of Adopted Area

Appendix C – Supplemental Figures

Data for simulations with no side payment mechanism come from 301 Monte Carlo sets, each including a 6x6x4x3 sweep of the policy variables, for a total of 130,032 modeling runs.

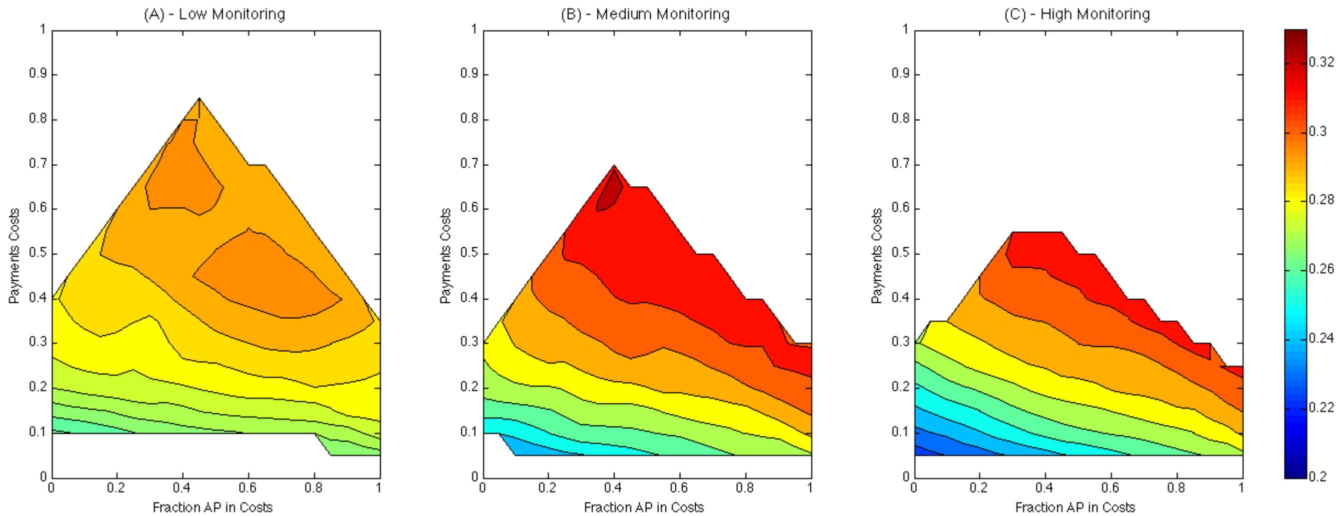

Figure A1: Contiguity of plots adopting conservation practice in simulations with no side payment mechanism, as function of total program spending on payments (X axis, along left axis within each panel) and the percentage of payments coming as agglomeration payments (Y axis, along right axis within each panel), for three different conditions of monitoring effectiveness (Low – per-area penalty of 1000 and chance of being caught 0.4; Medium – per-area penalty of 1500 and chance of being caught 0.6; High – per-area penalty of 2000 and chance of being caught 0.8). Surface is irregularly shaped as both x- and y-axis variables are modeled outcomes, not input variables. Spending on payments (X axis) is re-scaled to range from 0 to 1. Contiguity is calculated as the fraction of planted area in the neighborhood of an adopting plot, that is also adopting conservation agriculture (where neighborhood is the radius defined by the agglomeration payments program).

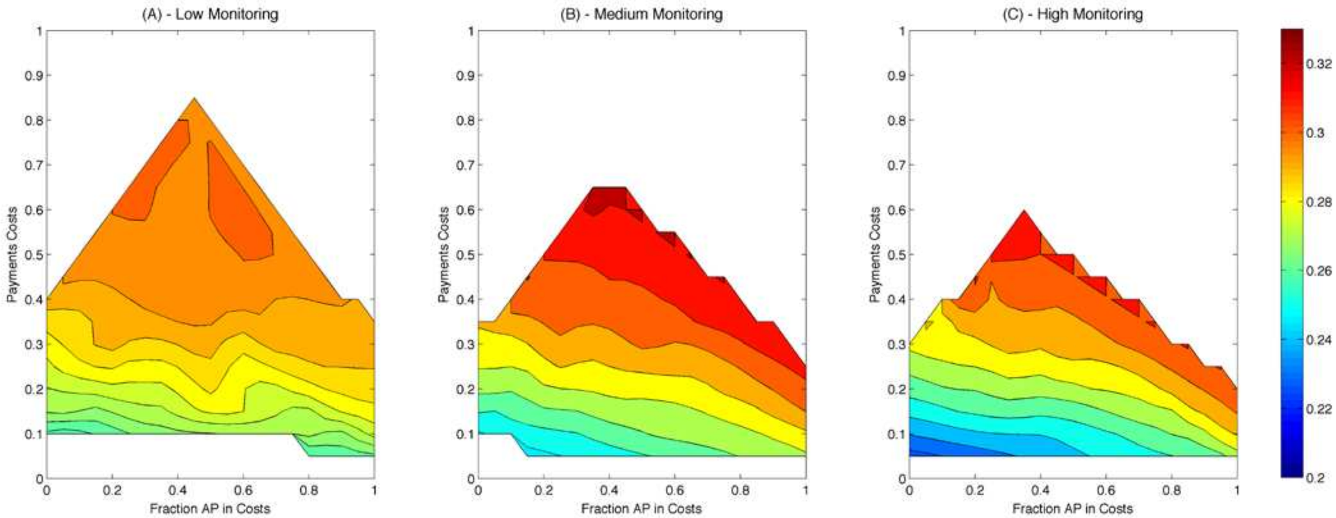

Figure A2: Proportion of total plot area adopting conservation practice in simulations with no side payment mechanism, as function of total program spending on payments (X axis, along left axis within each panel) and the percentage of payments coming as agglomeration payments (Y axis, along right axis within each panel), for three different conditions of monitoring effectiveness (Low – per-area penalty of 1000 and chance of being caught 0.4; Medium – per-area penalty of 1500 and chance of being caught 0.6; High – per-area penalty of 2000 and chance of being caught 0.8). Surface is irregularly shaped as both x- and y-axis variables are modeled outcomes, not input variables. Spending on payments (X axis) is re-scaled to range from 0 to 1.

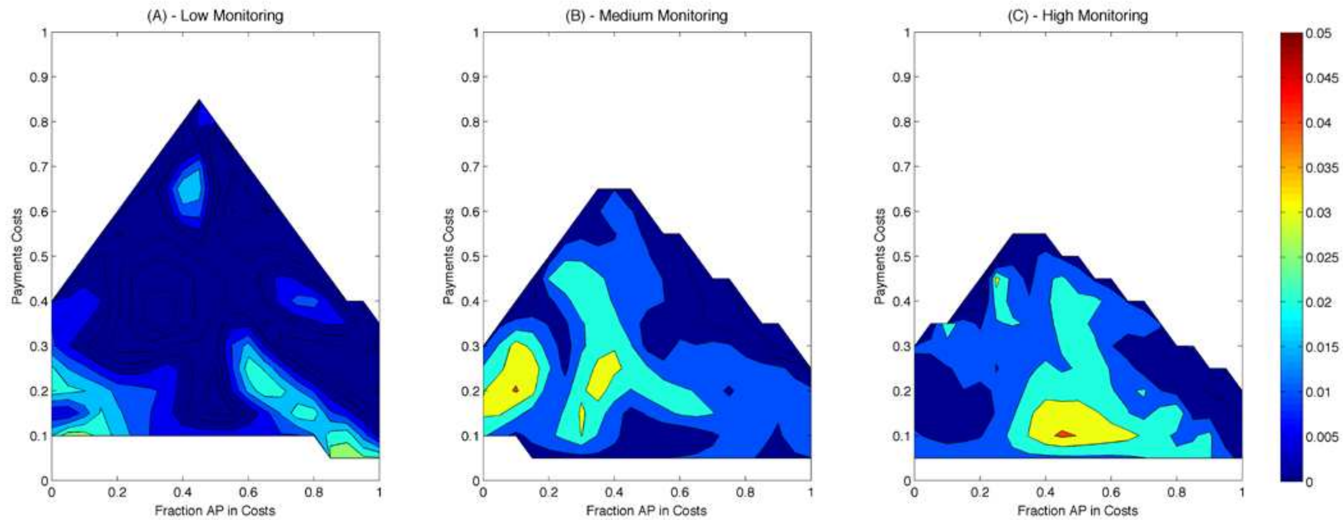

Figure A3: Relative fraction of proportion of total plot area adopting conservation practice attributable to side payments, as function of total program spending on payments (X axis, along left axis within each panel) and the percentage of payments coming as agglomeration payments (Y axis, along right axis within each panel), for three different conditions of monitoring effectiveness (Low – per-area penalty of 1000 and chance of being caught 0.4; Medium – per-area penalty of 1500 and chance of being caught 0.6; High – per-area penalty of 2000 and chance of being caught 0.8). Surface is irregularly shaped as both x- and y-axis variables are modeled outcomes, not input variables. Spending on payments (X axis) is re-scaled to range from 0 to 1. Relative fraction is calculated as $(A - A_{no\_side})/A_{no\_side}$ where A are the values from the adopted area surfaces in Figure 4, and A_{no_side} are the values from the adopted area surfaces in Figure A1 from the values in surfaces of Figure A1.

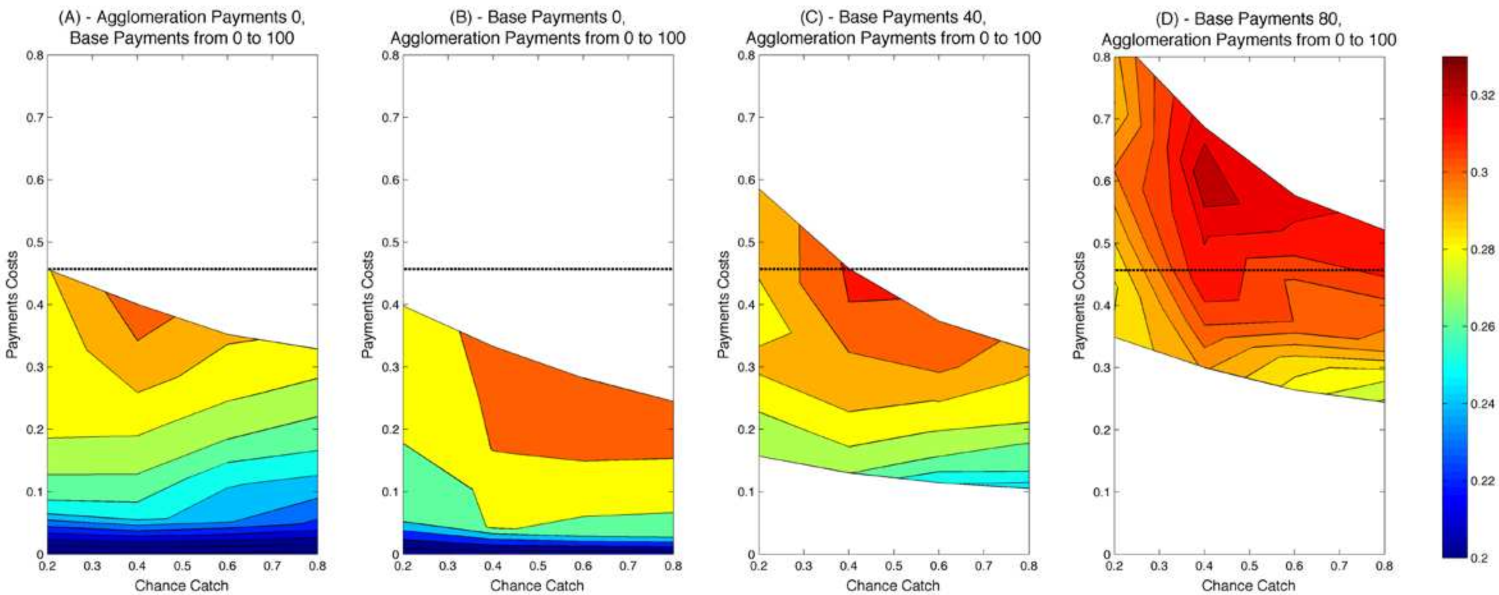

Figure 5: Proportion of total plot area under conservation practice in simulations with no side payment mechanism, as function of spending on payments and the likelihood of being caught, for several subsets of data: A) all model cases where agglomeration payments are 0 (with base payments spanning 0 to 100), B) all model cases where base payments are 0 (with agglomeration payments spanning 0 to 100), C) all model cases where base payments are 40, and D) all model cases where base payments are 80. Color (from blue through red) represents proportion of total area under conservation practice; contours of same color have same level of area under conservation practice. Black dashed line represents the maximum program spending in model cases with only base payments. The "L-shaped" contours in B-D (absent in A) indicate that different combinations of monitoring effort (proxied by the likelihood of being caught) and agglomeration payments may lead to the same adoption area. For example, the maximum area using base payments only (at a cost of 0.375 and monitoring effort of 0.4) in A, can be achieved with base and agglomeration payments at half the cost (~0.18) at the same monitoring effort in B, or at a higher cost with lower monitoring, or *vice versa*. Surface is irregularly shaped as y-axis variable is a modeled outcome, not an input variable.