

Identification of Manipulation in Receding Horizon Electricity Markets

Paul Scott and Sylvie Thiébaux

Abstract—Future distribution network markets will need to be flexible enough to enable the participation of small-scale customers with distributed energy resources. We propose using a receding horizon market that can manage the state-dependent decisions and large uncertainties of these participants. Unfortunately, this added flexibility creates new opportunities for agents to manipulate the actions of others by misrepresenting their true preferences for energy. This paper investigates this form of market manipulation in detail by first formalising the notions of receding horizon inconsistency and manipulation. We present a method for experimentally calculating the impact of a manipulative agent, and run it on two market settings, one based on a wholesale market and the other on a market providing distribution network support. We develop simple privacy-preserving indicators to identify inconsistency and manipulation, and demonstrate a difference in the behaviour of uncertain and manipulative agents. When using these indicators in a test for receding horizon manipulation, we correctly identify more than 95% of the days in which the greedy agents undertake the most harmful form of manipulation. Market operators can use these tools to run the system closer to its social optimum by restricting or penalising manipulative actions.

Index Terms—Demand Response, Receding Horizon, Manipulation, Multi-Agent Systems, Markets, Distributed Energy Resources

I. INTRODUCTION

WHOLESALE electricity markets are used throughout the world for the safe and efficient coordination of supply and demand. Network operators are starting to consider how similar markets, operating down at the distribution level, could help to manage capacity and voltage constraints as well as better align load with generation¹. Distributed energy resources (DERs), such as PV, battery storage, electric vehicles and smart appliances, are providing the necessary flexibility to make this possible. There is a great opportunity to provide new solutions to existing network problems, as well as to avoid future network problems that a large uptake of *uncoordinated* DERs can cause.

Crucially, any such distribution market needs to be designed to make it easy for residential and other small-scale electricity customers to participate. One aspect is the use of customer energy management systems (EMSs) to automate market interactions. Another is ensuring that the market can operate well despite the large uncertainties that small-scale customers

are exposed to. This paper investigates a market mechanism that operates over a receding horizon in order to provide the flexibility for participants to recover from uncertain events. It focuses on market manipulation in this market, because enhanced flexibility not only enables uncertain agents to participate, it also creates new opportunities for manipulation.

Existing wholesale markets help generators and large loads manage their uncertainties by offering a range of different markets operating at different time scales, including long-term futures, day-ahead, intra-day and frequency control markets. Day-ahead markets aid price discovery and enable participants to plan out their actions over a forward horizon — something that is particularly important when the generating equipment or loads under their control are state-based in nature (e.g., generator ramping rates and battery state of charge). Intra-day markets enable participants to recover from any uncertain events that occur in closer to real-time. This allows them to trade power that they would otherwise not be able to fully commit to in a day-ahead market.

In the European Power Exchange (EPEX) the day-ahead and intra-day markets are separate, while in the National Energy Market (NEM) [1] in Australia they have been integrated. In the NEM, generators submit day-ahead offers, but they may rebid throughout the trading day to adjust their offered capacity. The market is cleared in close to real-time (every 5 minutes) instead of committing generation to a future dispatch. Compared to wholesale markets, participants in a distribution market are exposed to even higher degrees of operational uncertainty due to their small size. For this purpose we propose a market that operates over a receding horizon, which is a generalisation of the NEM approach to provide more flexibility. Receding horizon approaches have been suggested and experimented with in closely related contexts including generator dispatch control, real-time price demand response and distributed network management [3], [30], [31].

Figure 1 shows the structure of the receding horizon approach, where a horizon consists of a number of time steps spanning a forward period of time (e.g., 24 hours). Participants provide their bids for the horizon and the market is cleared, but agents are only required to commit to the outcome of a horizon's first time period. The problem then moves onto the next horizon and the process repeats. This not only makes it possible for smaller consumers to participate, it also provides a better outcome because the dispatch can be reoptimised in near-to real-time².

Paul and Sylvie are in the College of Engineering and Computer Science at the Australian National University, Canberra, Australia.

¹In addition to ongoing industry-partnered research projects on this topic, the Australian Energy Market Commission (AEMC) is currently undergoing a project to explore possible market models <http://www.aemc.gov.au/Markets-Reviews-Advice/Distribution-Market-Model>.

²This real-time decision making and model-based approach has made receding horizon (model predictive control) techniques widely successful in the field of control [2].




Figure 1: A receding horizon market considers the exchange of power over a forward horizon of T time steps, which recedes as the first time step is reached in real-time. The market clears each horizon, but agents only need to commit to the outcome of the first time step in each horizon.

Such markets that allow rebidding throughout the day enable new market manipulation opportunities above and beyond what is possible in a one-shot day-ahead market. A receding horizon approach provides agents with a great amount of flexibility to recover from uncertain events, but with this comes a range of opportunities for agents to manipulate prices for their own favour; normally at the expense of the social optimum. In a typical scenario, a strategic agent exaggerates how much energy they require in a future time step in order to temporarily inflate the price. Flexible agents will then commit to consuming earlier, which will lower the demand and price for power in the later time step once the strategic agent reverts to their true requirements. We consider this illegitimate market behaviour, but the challenge is that such behaviour can in general be indistinguishable from the legitimate actions of agents that are uncertain about their energy requirements.

This paper investigates how much of a problem manipulation of a receding horizon electricity market is, and what can be done about it. It focuses on a class of receding horizon market mechanisms that pay/charge participants at a marginal clearing price, which are characterised as producing Karush-Kuhn-Tucker solutions when clearing the market (sections II-A and III). In general these mechanisms are not incentive compatible, but there exist mechanisms in this class that are computationally efficient, budget balanced, and, assuming convex preferences, produce socially optimal outcomes *when* agents act truthfully. These mechanisms can make use of distributed optimisation techniques, enabling them to scale for use in distribution markets with many participants [3]–[7].

Our contributions are threefold, starting with a thorough formalisation of receding horizon inconsistency and manipulation (section IV) for this class of mechanisms.

Our second contribution is the presentation of a method for calculating an optimal strategy for a manipulative agent in a complete information setting (section V). This agent is used to empirically investigate how much a single agent can gain by manipulating the mechanism, and how this harms the social objective (section VI). In the market settings we investigate, we find that when agents form a coalition with a large enough share of the market energy, they can gain an advantage by manipulation. However, in practice incomplete information and a high computational burden will make it challenging for them to obtain the full benefits demonstrated here.

Our third contribution is the development of privacy-preserving indicators that can test whether an agent is ma-

nipulative when we limit the uncertainty that agents exhibit (section VII). Supporting theoretical and empirical results show how this approach can identify agents that are driving change, and distinguish a manipulative agent from an uncertain agent (section VIII). This provides the full benefits of a receding horizon mechanism whilst significantly limiting the opportunities for manipulation.

The final sections of this paper discuss related work in light of our contributions before concluding.

II. MARKET PROBLEM

The goal of electricity markets is to facilitate the safe and efficient supply and consumption of electricity. For our problem this translates to ensuring supply remains in balance with demand at all moments³, and that the outcome maximises the social welfare of participants. If we assume we have full access to the power usage preferences of all participants (ignoring the complicating agent interactions with the market), we can define the desired market outcome as a constrained optimisation problem, which we refer to as the *power balancing problem*.

This section explains the power balancing problem and some simple market participant models that will be used in our experiments. The section after this will bring back the agent interactions by introducing a class of market mechanisms that attempt to solve this power balancing problem in a competitive market environment.

A. Power Balancing Problem

Consider a market with A participants and a horizon that consists of T equally sized time steps, e.g., 24 time steps each an hour in duration to make up a 24 hour horizon. Each agent $i \in \{1, \dots, A\}$ has for each horizon $h \in \{1, 2, \dots\}$ a power profile⁴ $P_{h,i} \in \mathbb{R}^T$, a cost function $f_{h,i} : \mathbb{R}^T \mapsto \mathbb{R}$ and N_i constraint functions $\forall j \in \{1, \dots, N_i\} : g_{h,i,j} : \mathbb{R}^T \mapsto \mathbb{R}$. We restrict these functions to be continuously differentiable or convex, which are conditions of the class of market mechanisms we introduce in the next section.

An agent's cost and constraint functions encode their preference for consuming/producing electricity over the horizon of interest. Importantly these can change between horizons, both in order to account for changes in their forecasts and because each new horizon drops the oldest time step and picks up a new one (see Figure 1).

A common interpretation of the overall social welfare is as the sum of each agent's utility function (in our problem the negative sum of all agent cost functions). For a given horizon h , the power balancing problem is then:

$$\min_{P_{h,i}} \sum_{i=1}^A f_{h,i}(P_{h,i}) \quad (1)$$

$$\text{s.t. } g_{h,i,j}(P_{h,i}) \leq 0 \quad \forall i \in \{1, \dots, A\}, j \in \{1, \dots, N_i\} \quad (2)$$

$$\sum_{i=1}^A P_{h,i} = 0 \quad (3)$$

³More sophisticated techniques can take this one step further and model safe network power flows, which we discuss as an extension in section X.

⁴Using load convention where positive is consumption, negative production.

where the power profiles of each agent across the horizon form the decision variables, and equation (3) ensures supply and demand is balanced in each time step. Our definitions and theoretical results are formulated for this problem, while a modified terminating receding horizon (TRH) version is used in experiments to reduce the problem to a finite number of horizons, enabling a more straight-forward way of comparing results between experiments.

The TRH version consists of T consecutive horizons $h \in \{1, \dots, T\}$, which truncate at time step T . We model this as T perfectly overlapping horizons, each of length T , and force some of the variables from the later horizons to take on values from earlier horizons (as illustrated in Figure 2). This is achieved by adding the following constraint to (1–3):

$$P_{h,i,t} = P_{h-1,i,t} \quad \forall i \in \{1, \dots, A\}, t < h \quad (4)$$




Figure 2: Representation of the TRH horizons. Shaded areas represent the time steps in which variables are constrained to take on values from the previous horizon (according to equation (4)).

In this formulation the cost and constraint functions of each agent now contain the exact same time steps. This means that we would only expect these functions to change between horizons if the agent is uncertain and gains more information over time or, as we will discuss later, if they are intentionally attempting to manipulate the mechanism.

B. Agents

We present four simple agents that are used in our experiments. They provide the basic important behaviour of a large range of other network agents without making the strategic agent formulation overly complicated to implement or solve.

1) *Generator*: The generator has a quadratic cost function $f(P) = \sum_{t=1}^T (\frac{1}{2}\Psi_t P_t^2 + \psi_t P_t)$, where $\Psi \in \mathbb{R}_{\geq 0}^T$ and $\psi \in \mathbb{R}^T$, which approximates the increasing marginal cost of power production typically observed in wholesale electricity markets. The generator also has power limits such that for each time step $P_t \in [P, \bar{P}]$.⁵

2) *Fixed Load*: This agent has no cost function or decision variables, it simply has a fixed power profile which it must consume. This represents a network load that has no flexibility.

3) *Battery*: We use a very simple battery model that has no losses or deterioration. It has a zero cost function, a positive capacity $\bar{E} \in \mathbb{R}_{>0}$, a starting charge $E_0 \in [0, \bar{E}]$, and charge and discharge power limits $P_t \in [P, \bar{P}]$. The battery has a series of constraints that ensures that the state of charge remains within its limits at each time step: $0 \leq E_0 + \sum_{\tau=1}^t P_\tau \leq \bar{E}$.

⁵These limits can be represented in the general agent formulation as a series of constraint functions, e.g., $g_1(P) = P - P_1$, $g_2(P) = P_1 - \bar{P}$, etc.

4) *Deferrable Load*: Deferrable loads have no cost function but they have constraints that force them to consume a total amount of energy $E = \sum_{t=1}^T P_t$ over the horizon, and time dependent constraints on when it can consume this energy $P_t \in [0, \bar{P}_t]$. Deferrable loads capture the state-based nature of smart appliances whose load can be deferred in time including dish washers and hot water heaters.

III. MARKET MECHANISM

A market mechanism defines the rules for how participants can bid into a market, how the market is cleared and then who gets paid what. This section goes into some of the details of our proposed mechanism which we simply refer to as “the RH mechanism” or “our mechanism” throughout this paper.

For a particular horizon, agents bid into the market by supplying the mechanism with convex cost and constraint functions that represent their preferences for power over the horizon. The mechanism then clears the market by finding an optimal solution to the resulting power balancing problem. The solution consists of a power profile for each agent $P_{h,i}$ and the marginal price of electricity $\lambda_h \in \mathbb{R}^T$. Agents are then dispatched at this resulting power and charged/paid at the marginal price for the first time step in the horizon. This process then repeats for the next horizon.

A. KKT Conditions

We use Karush-Kuhn-Tucker (KKT) points to mathematically capture an optimal solution to the power balancing problems that are solved for each horizon. Under certain regularity conditions (constraint qualifications) on the problem functions, a KKT point becomes necessary for optimality [8], [9]. Under other conditions such as convexity, a KKT point is sufficient for optimality [10].

A KKT point for a convex problem (an optimal solution) can be efficiently computed, and this computation can be distributed across many nodes using techniques such as ADMM [11]. In practice this can speed up the computation and preserves a greater degree of agent privacy, since distributed optimisation techniques do not need to have direct access to an agent’s cost and constraint functions. This scalability is an important consideration in a power systems context, where there are potentially many thousands of participants with different preferences for energy over the horizon.

When the cost and constraint functions are continuously differentiable, the KKT conditions for an agent i in horizon h are:

$$\nabla f_{h,i}(P_{h,i}) + \sum_{j=1}^{N_i} \mu_{h,i,j} \nabla g_{h,i,j}(P_{h,i}) + \lambda_h = 0 \quad (5)$$

$$g_{h,i,j}(P_{h,i}) \leq 0 \quad \forall j \in \{1, \dots, N_i\} \quad (6)$$

$$\mu_{h,i,j} g_{h,i,j}(P_{h,i}) = 0 \quad \forall j \in \{1, \dots, N_i\} \quad (7)$$

$$\mu_{h,i,j} \geq 0 \quad \forall j \in \{1, \dots, N_i\} \quad (8)$$

where $\mu_{h,i,j} \in \mathbb{R}$ and $\lambda_h \in \mathbb{R}^T$ are the KKT multipliers for the agent’s constraints and the power conservation constraint (3) respectively. These conditions are commonly referred to as stationarity (5), primal feasibility (6), complimentary slackness

(7) and dual feasibility (8). Particular powers $P_{h,i}$ and prices λ_h satisfy the KKT conditions for agent i if there exists some $\mu_{h,i,j}$ where equations (5–8) hold.

The KKT conditions for the overall problem in horizon h is the combination of KKT conditions for all agents along with the power conservation constraint (3).

When dealing with convex functions, which might be non-differentiable, we allow a subgradient to be used in place of the gradient. Therefore, if f is a convex function, then we define $\nabla f(P) \in \partial f(P)$, where $\partial f(P)$ is the subdifferential (set of subgradients) of f at P . Using this notation the KKT stationarity condition (5) remains the same.

B. Payments

In this mechanism agents are paid (charged) for their power production (consumption) at the marginal prices $\lambda_h \in \mathbb{R}^T$. The expected cost $c_{h,i}$ for an agent i in horizon h is the combination of their cost function and expected payments:

$$c_{h,i} := f_{h,i}(P_{h,i}) + \lambda_h^\top P_{h,i} \quad (9)$$

which assumes that agents have quasilinear utilities and are risk neutral. This is the *expected* cost when solving for horizon h , because only the first time step in each horizon is acted on, while the remaining time steps that overlap with future horizons are hypothetical. In the TRH version of the problem, the costs associated with the last horizon represent the actual costs exhibited by agents because the full market problem has been cleared and committed to at this point (see Figure 2).

Marginal pricing is used in a number of wholesale electricity markets including in Australia, New Zealand and various regions of the US. When they are extended to account for physical power flows across the network they are referred to as locational marginal prices.

C. Limits

1) *Power Limits*: In order to participate in the mechanism, each agent must first negotiate a contract that restricts how much power they can consume and supply. This contract establishes limits such that $\forall t \in \{1, \dots, T\} : P_{i,t} \in [\underline{P}_i, \bar{P}_i]$. They can be set at the physical limits of the equipment connecting the agent to the rest of the network, or to a subinterval of these physical limits if the agent does not need the extra capacity.

2) *Price Limits*: A second limit is on the allowed prices λ , which is commonly referred to as a market cap in wholesale electricity markets. This can be set at a point where participants become indifferent towards being disconnected from the network (i.e. the point where load shedding takes place). If the solver finds a KKT point with costs outside the price caps or if no KKT point is found at all, then the market operator can shed loads in an attempt to find a feasible solution.

IV. RECEDING HORIZON MANIPULATION

Market manipulation is the deliberate act of artificially inflating or deflating market prices. Mechanisms that are incentive compatible are designed in a way so that agents

cannot gain any advantage by manipulating the market. Mechanisms such as VCG jointly achieve incentive compatibility and social optimality; however, VCG only achieves this under strict conditions (e.g., no coalitions) and has other drawbacks including computational complexity and a general lack of budget balance [12]. Furthermore, VCG does not flexibly allow the participation of uncertain agents.

The RH mechanism in contrast provides uncertain agents with the flexibility to participate, but it is not incentive compatible. Agents can manipulate the mechanism by misrepresenting their true preferences (their “type”) to the mechanism, which we refer to as acting *untruthfully*. Fortunately, as explored in section VII, in a practical setting there are some simple ways to deter the more extreme acts of manipulation without significantly impacting uncertain agents.

This section provides an example and formal definition of receding horizon manipulation of our mechanism, but first some notation is introduced. When we focus on an arbitrary pair of consecutive horizons h and $h + 1$, we drop the horizon subscript from functions, powers and prices, and instead indicate the later horizon with a prime (e.g., $f_i := f_{h,i}$ and $f'_i := f_{h+1,i}$). To make it easier to directly compare consecutive horizons, we *time shift* powers and prices from the later horizon so that they line up with the values from the earlier horizon (see Figure 3). These vectors are marked with an asterisk, and for power (same for prices) it is defined as:

$$P_{i,t}^* := \begin{cases} P_{h,i,t} & \text{if } t = h \\ P_{h+1,i,t} & \text{if } h < t < h + T \end{cases} \quad (10)$$



Figure 3: For consecutive horizons, the shaded time steps represent the values used for the time-shifted price and power vectors.

We define vectors for the change in power and price as $\Delta P_i := P_i^* - P_i$ and $\Delta \lambda := \lambda^* - \lambda$. We also define the change in cost with respect to the earlier horizon as $\Delta c_i := f_i(P_i^*) + \lambda^{*\top} P_i^* - f_i(P_i) - \lambda^\top P_i$.

A. Truthfulness and Uncertainty

We use a hat to indicate an agent’s best estimate of their cost and constraint functions for a particular horizon, e.g., $\hat{f}_{h,i}$ and $\hat{g}_{h,i,j}$. If an agent is uncertain these might not accurately represent the preferences that will eventuate for the agent later on, but it is their best estimate at that moment in time. The functions that an agent uses to interact with the mechanism are $f_{h,i}$ and $g_{h,i,j}$, which they are free to choose irrespective of their actual preferences. When these are exactly equal to their best estimates, the agent is truthful. Formally:

Definition 1 (Truthfulness). For a horizon h , let $f_{h,i}$ and $g_{h,i,j}$ be the functions that agent i uses with the mechanism, and $\hat{f}_{h,i}$ and $\hat{g}_{h,i,j}$ be the functions the agent believes are most likely true at that moment in time. Agent i is *truthful* in a horizon h if $f_{h,i} = \hat{f}_{h,i}$ and for all $j : g_{h,i,j} = \hat{g}_{h,i,j}$. Otherwise agent i is *untruthful*.

Without a receding horizon a one-shot day-ahead market is still susceptible to manipulation, but this form of manipulation is generally well studied and understood. For this reason, we focus our receding horizon manipulation definition on the form of manipulation that is unique to the receding horizon part of the mechanism. Before delivering this definition we provide an example to highlight this form of manipulation.

B. RH Manipulation Example

Consider a two time step TRH market problem with a generator, deferrable load and fixed load. Assuming no uncertainty and that the generator and deferrable load act truthfully, the fixed load can gain an advantage in certain situations. Assume the generator has a quadratic cost term $\hat{\Psi} = 1$ for both time steps, the deferrable load must consume $\hat{E} = 2$ units of power (it does not care when) and the fixed load needs to consume more power in the second interval than the first $\hat{P} = (0.5, 1.5)$. By lying in just the first horizon (e.g., by saying it needs $P = (0.5, 2.5)$), the fixed agent can temporarily inflate the second time step price, causing the mechanism to fully dispatch the deferrable load in the first time step.

Table I compares the outcome with and without the fixed load manipulating. The results show how the total system costs increase and the fixed load agent’s costs decrease when it is manipulating.

Table I: Manipulation example comparison.

	P Gen	λ	Total Cost	Agent Cost
Manipulating	(2.5, 1.5)	(2.5, 1.5)	4.25	3.5
Truthful	(2.0, 2.0)	(2.0, 2.0)	4.00	4.0

C. RH Consistency and Manipulation

We consider an agent to have performed receding horizon manipulation when two criteria are met: they are untruthful, and their change in behaviour between horizons has a tangible affect on the market. An agent can be untruthful and have no impact, or they could be truthful and have an impact (e.g., if they are uncertain); neither case is considered as receding horizon manipulation. The following notion of receding horizon consistency captures whether or not an agent has had such impact on the RH mechanism.

An agent is consistent between consecutive horizons when the preferences they provide in the earlier horizon could have produced the result obtained in the later horizon (for those time steps that overlap). Formally:

Definition 2 (Receding Horizon Consistency). Let the KKT points for the power balancing problem in two consecutive horizons be (P_i, λ) and (P'_i, λ') for agent i . If the later horizon time-shifted point (P_i^*, λ^*) satisfies the earlier horizon KKT conditions (5–8) for agent i , then agent i is *consistent* between the horizons. Otherwise agent i is *inconsistent* between the horizons.

It is possible for an agent to change their preferences between horizons and be consistent, as this definition focuses on the outcomes by asking whether an agent’s earlier functions

could produce the same outcome in the later horizon. If so then they have not changed their functions in a way that has altered the outcome in a consequential way.

An agent manipulates the receding horizon part of the RH mechanism when it is untruthful in order to create an inconsistent result. Uncertain agents may be inconsistent, but if they are acting truthfully then they are not manipulating the receding horizon mechanism. Formally:

Definition 3 (Receding Horizon Manipulation). If agent i is inconsistent between consecutive horizons and untruthful in the earlier horizon, then agent i is *manipulating* the receding horizon problem between these horizons.

V. GREEDY AGENT STRATEGY

In general, the RH mechanism results in an incomplete information game: agents might have some beliefs about the preferences or “types” of other agents, but they typically do not know these exactly. A conventional game theory approach is to look for market equilibria where all agents are acting strategically with a set of beliefs about other agent types. This analysis could identify certain information poor settings where it is impossible for an agent to manipulate the system in expectation.

While we hope to undertake such a detailed analysis in future work, for now we focus on a specific concrete case where we investigate the maximum benefit a single manipulative agent (or coalition of agents) could achieve, and what harm this causes to the system. We assume this “greedy” agent has complete information about all other agents, and that the rest act truthfully. This simplified setting is useful because it captures the best-case benefit a manipulative agent could hope to gain from the system. It also provides a test case for our manipulation identifiers, which tease apart manipulative agents from uncertain agents.

This section describes the approach we take for calculating an optimal strategy for a greedy agent. The approach involves solving a bilevel program, where the greedy agent at the top level chooses its optimal preferences given that it knows how the mechanism will respond at the lower level [13]. In general, even linear bilevel programming is NP-hard [13, chap. 6]. This will dramatically limit the size and type of problems that can be solved optimally. We will use bounds and solve simple instances in order to gain insights despite the complexity.

The TRH problem is used in the formulation with the greedy agent having an index $i = 1$. The greedy agent chooses its power for each horizon, and there is no uncertainty: for all h let $\hat{f}_{h,i} = \hat{f}_i$ and $\hat{g}_{h,i,j} = \hat{g}_{i,j}$. We assume that there are market cap prices λ, λ' and that the greedy agent will avoid a solution which violates these because of the risk of being disconnected from the network.

We make the *optimistic* bilevel assumption, which allows the greedy agent to choose between lower level solutions if more than one exists for a given upper level decision. This allows the problem to be immediately flattened to a single

level program:

$$\min_{P_{h,i}, \lambda_h} \hat{f}_1(P_{T,1}) + \lambda_T^T P_{T,1} \quad (11)$$

$$\text{s.t. } \hat{g}_{1,j}(P_{T,1}) \leq 0 \quad \forall j \in \{1, \dots, N_1\} \quad (12)$$

$$P_{h,i,t} = P_{h-1,i,t} \quad \forall i \in \{1, \dots, A\}, \quad (13)$$

$$h \in \{2, \dots, T\}, t < h$$

$$\lambda_{h,t} = \lambda_{h-1,t} \quad \forall h \in \{2, \dots, T\}, t < h \quad (14)$$

$$P_{h,1} \in [P_1, \bar{P}_1]^T \quad \forall h \in \{1, \dots, T\} \quad (15)$$

$$\lambda_h \in [\lambda, \bar{\lambda}]^T \quad \forall h \in \{1, \dots, T\} \quad (16)$$

$$(P_{h,1}, \dots, P_{h,A}, \lambda_h) \in \text{KKT}_h \quad (17)$$

$$\forall h \in \{1, \dots, T\}$$

The objective is to minimise the greedy agent's final costs, which are given by the values in the last horizon: $h = T$. The greedy agent's true constraints are enforced on the final powers with inequality (12). Equations (13) and (14) tie together powers and prices between horizons (those that have been finalised), in accordance with the TRH formulation. Constraints (15) and (16) bound the greedy agent's power and the prices according to the limits discussed in section III-C.

The lower level problem is the solution of the market mechanism (the power balancing problem) in each horizon, given the greedy agent's untruthful preferences. Without loss of generality to what the greedy agent can achieve, these preferences are provided to the mechanism as fixed powers in each horizon. Representing the solution of the lower level problem by its KKT conditions (17) has allowed the bilevel program to be flattened into a single level. KKT_h is defined to be the set of feasible KKT points for the power balancing problem in horizon h with the greedy agent's powers fixed. The resulting stationarity (5) and complementary slackness (7) equalities from the KKT conditions, and the bilinear cost term in the objective (11) are all possible sources of non-convexity.

The KKT conditions are composed of those associated with each truthful agent in the problem. In the next section we show how we formulate these conditions for a truthful generator. A similar approach was used for the battery and deferrable loads (the fixed load is trivial), with further details and derivations provided in the PhD thesis [14, chap. 5].

A. Generator KKT Conditions

The KKT conditions (which contain bilinear terms) are transformed into a series of mixed-integer linear constraints, so that more established methods for global optimisation can be applied. Binary variables $z_{h,i,t}^u$ and $z_{h,i,t}^l$ are introduced for the power upper and lower bounds, and the KKT multipliers for these bounds can be combined into a new variable $\nu_{h,i,t} = \mu_{h,i,t}^u - \mu_{h,i,t}^l$. Big-M style constant bounds $\nu_{i,t}$ and $\bar{\nu}_{i,t}$ are used for these new multipliers, The reformulation is $\forall t \geq h$:

$$\Psi_{i,t} P_{h,i,t} + \psi_{i,t} + \nu_{h,i,t} + \lambda_{h,t} = 0 \quad (18)$$

$$\nu_{h,i,t} \leq z_{h,i,t}^u \bar{\nu}_{i,t}, \quad P_{h,i,t} \geq z_{h,i,t}^u (\bar{P}_i - P_i) + P_i \quad (19)$$

$$\nu_{h,i,t} \geq z_{h,i,t}^l \nu_{i,t}, \quad P_{h,i,t} \leq z_{h,i,t}^l (P_i - \bar{P}_i) + \bar{P}_i \quad (20)$$

These constraints represent the stationarity, dual feasibility, primal feasibility and complementary slackness conditions of the original formulation. The market price caps and the stationarity condition provide bounds for the multiplier $\nu_{h,i,t}$:

$$\nu_{i,t} = -\bar{\lambda} - \Psi_{i,t} \bar{P}_i - \psi_{i,t} \quad (21)$$

$$\bar{\nu}_{i,t} = -\lambda - \Psi_{i,t} P_i - \psi_{i,t} \quad (22)$$

VI. GREEDY AGENT EXPERIMENTS

We develop two market settings inspired by real power networks, to investigate how much a greedy agent has to gain by manipulating, and how much this hurts the overall social outcome. Section VIII uses these market settings to test an approach to identifying receding horizon manipulation.

A. Wholesale Market Setting

The first market setting is inspired by existing wholesale electricity markets. Spot wholesale price and demand data was obtained for the New South Wales region of the Australian NEM for the trading day starting on the 23rd of February 2017, which experiences a peak load of 11 GW. Instead of individually modelling each generator in the system, we use a single generator tuned to produce the same overall trend in prices: an increasing marginal cost of production. We performed linear regression on this data and obtained coefficients of $\Psi = 0.0114$ \$/MWh and $\psi = -15.06$ \$/MWh for our aggregate generator (constant over time).

We use two fixed load agents to represent the demand, one greedy and one truthful. We parameterise the split of the total demand between these two agents. Finally, we add some flexible DERs into the network in the form of batteries. By performing some simple experiments, we found that we could get equivalent results by representing many small batteries with a single large battery, so we went with a single battery who's capacity and power rating are parameterised.

B. Distribution Market Setting

The second market setting is for a market on a distribution network designed to coordinate residential batteries to manage network congestion. In this problem, a conductor in a feeder reaches its capacity during peak events. The total peak downstream load is around 1300 kW and the conductor limit is 950 kW. There are two ways of managing this constraint, either dispatch a downstream diesel generator or discharge downstream residential battery systems to reduce the load during peak events. This is a real network problem that is being investigated in an ongoing trial with a partner DNSP.

There are 750 houses downstream of the network constraint, which we split up into three groups. The first is a group of houses that work together as a collective to manipulate the market which we model as a greedy fixed load. The second is a group of houses that act truthfully modelled as a fixed load. The third is a group of houses with batteries that also act truthfully, approximated as a single large battery.

The diesel generator has a price of 1 \$/kWh ($\psi = -1$). A second generator model is used to represent the power that enters the constrained region of the feeder from the constrained

conductor. It has zero cost and a capacity limit equal to the conductor capacity.

C. Problem Instances

A number of problem instances were produced from these market settings by choosing particular parameter values and randomising the loads of the greedy and truthful fixed loads. TRH problem instances were produced with horizons of length $T = 24$ (24 hours). Calculating optimal agent strategies in this way becomes intractable with large numbers of agents and horizons with many time steps. In order to achieve a tractable problem, we restrict ourselves to a small number of agents which aggregate together the behaviour of many other smaller agents. This provides a good approximation for the agents we consider, which is sufficient for our intentions of gauging the potential scale of the effects from receding horizon manipulation and demonstrating the ability of our identification techniques.

Only a single agent is acting strategically in each case. All other agents act truthfully and there is no uncertainty in the problem for the first set of experiments.

When the greedy agent is a fixed load the problem can be reduced to a MILP. These problems were solved using Gurobi 7. After presolve, the greedy agent problem instances for the wholesale market setting have around 1000 continuous and 1000 binary variables. A solving cutoff time of 10 minutes was used for each instance. The resulting optimality gap was recorded for the candidate solution at the end of this time.

D. Greedy Agent Impact

We explore the maximum impact a strategic greedy agent (modelling a cooperative of smaller greedy agents) can have in the wholesale and distribution market settings. By modifying the parameters in the problem we can explore how these impacts change with the relative proportion of greedy agents and flexible DERs. We generate 20 randomised days for each parameter setting and compare the mean change in the total system costs and the greedy agent's costs relative to the socially optimal outcome (all agents truthful so no manipulation) across the days.

Figure 4 shows how these cost changes vary for a combination of two different parameters. The first is the relative battery capacity/power output. This is measured as the total maximum battery power output \bar{P} as a percentage of the total network demand. The battery capacity \bar{E} (in kWh) is kept twice as large as the maximum battery discharge power (in kW), ensuring a full battery can discharge at its maximum rate for 2 hours. This percentage represents the level of battery uptake in the network, and the level of flexible DERs in the network.

As the number of flexible batteries increases (the three different colours), we see that the greedy agent has more opportunity to reduce their costs through manipulation (the dashed lines), and these actions also do more harm to the overall costs (the solid lines). Similarly there seems to be an overall trend that the greater percentage of the overall load that the greedy agent controls, the greater the manipulation payoffs,




Figure 4: Change in total cost (solid lines) and greedy agent cost (dashed lines) relative to social solution in the wholesale market setting. These are compared for different percentages of network battery uptake (total battery output as a percentage of total demand), and for different sized greedy agents (as a share of total load).

but only if there is enough flexible DERs in the network for it to influence.

A number of instances had a non-negligible optimality gap at the end of the 10 minute cutoff. The mean optimality gaps are displayed as error bars on the agent cost change curves as an indication of how much the candidate solutions might underestimate the true benefits that a greedy agent can gain. The worst gaps were experienced when the greedy agent has 10% of the total load and there is a 30% uptake in batteries. In this case, the error bar indicates that the average cost reduction for the greedy agent is somewhere in the range 1–2%. This was deemed sufficient for the purposes of this investigation, but if required a longer solve time could be employed to further reduce the gap. When the greedy agent has a larger share of the total load, and in the next set of experiments on the distribution market setting, the gaps are insignificant.

Figure 5 shows the same results but for the distribution market setting. We see that the greedy agent has the potential to achieve relative cost reductions that are much more significant in this setting, but it is also more sensitive to the particular setup of the problem. Here having a large share of the overall load is often a disadvantage, and too few or too many batteries in the network reduces the potential for manipulation. It is only in the intermediate case (enough flexible batteries for the greedy agent to influence, but not enough to completely solve the congestion problem) that there is opportunity for receding horizon manipulation.

Figure 6 shows an example solution for the greedy strategic agent in the wholesale market setting. In order to distinguish the solutions from the different horizons, we gradually increase the opacity of the lines representing power as the horizons progress. The greedy agent has a considerable impact on the battery and generator behaviour over time.

These results provide a bound on what would actually be achievable in a more realistic setting where there is incomplete information and where the problem scale creates computational challenges.




Figure 5: Change in total cost (solid lines) and greedy agent cost (dashed lines) relative to social solution in the distribution market setting. These are compared for different percentages of network battery uptake (percentage of households with batteries), and for different sized greedy agents (as a share of total load).




Figure 6: Solutions for horizons in a greedy agent instance in the wholesale market setting (greedy agent with 30% of total load, and battery maximum power at 30% of total load).

VII. RH MANIPULATION IDENTIFICATION

So far we have seen the potential impact that receding horizon manipulation can have in our markets settings. We now consider the position of a market operator who would like to reduce this. The approach we take is to develop indicators that can be used to test whether an agent is likely to be performing RH manipulating in a practical setting. Penalties can then be applied to the agents identified in this way, to deter future manipulation. We only assume we know the power profiles of each agent and prices in each horizon that the mechanism calculates. We do not have access to an agent’s beliefs about their true preferences or even the preferences they use to interact with the mechanism (as discussed certain distributed market clearing techniques do not have/need direct access to these preferences).

In general it is not possible to distinguish a manipulative agent from one who is just inconsistent due to their uncertainty, because we do not have access to an agent’s inner workings to determine if they are being truthful or

not. However, we find that if we restrict the ways in which an agent can be uncertain (the forms of uncertainty that the mechanism allows), then we can likewise restrict the undetectable strategies available to a manipulative agent. We develop simple to calculate indicators that can be applied to an agent’s power profiles for each horizon to test whether or not an agent is operating outside this acceptable range of behaviour.

One method for developing such tests, which we demonstrate in section VIII-D, is to run simulations of agents with acceptable forms of uncertainty in the market, and to use these results to establish appropriate parameters for the tests. Any behaviour outside these acceptable ranges would be an indication of manipulative behaviour. No test will perform perfectly in every case, and they will likely have to be developed over time to match the type of participants and the form of their uncertainty.

In a practical setting these tests will have some false positives and false negatives, because some agents will have uncertainties that lie outside the acceptable range, and other agents will still manage to perform some limited form of manipulation within the acceptable range. The market operator will need to tune these tests to strike the right balance of not penalising too many truthful agents, but also deterring the most significant forms of manipulation.

The first indicator we develop measures the degree of inconsistency between horizons. This can also be used to focus our attention on those horizons and those agents that are driving change in the market.

A. Revealed Preference Indicator

We derive a condition called revealed preference (RP) consistency that all receding horizon consistent agents will satisfy. When this condition is not satisfied, we know for certain that the agent is receding horizon inconsistent. This theory leads to the RP indicator which is a measure of how inconsistent an agent is.

RP consistency is similar to the RP activity rule proposed by Ausubel et al. [15]. Activity rules are used in iterative mechanisms, e.g., clock auctions and simultaneous ascending auctions, to limit the bidding strategies of agents [16, chap. 11]. Chapman and Verbic [17] used a RP activity rule within a single horizon to prevent agents from manipulating their bids in an electricity demand allocating clock auction.

The indicator we derive is similar to the RP activity rule, but it instead works with our notion of receding horizon consistency, and instead of enforcing a particular behaviour, we use it after the fact to get an indication of how inconsistent an agent is being.

Theorem 1 (Revealed Preference Consistency). *If an agent i is convex and consistent between horizons, then $\Delta\lambda^\top\Delta P_i \leq 0$.*

Proof. Assume i has convex cost and constraint functions and is consistent between consecutive horizons. Definition 2 implies that both the earlier horizon solution and later horizon time-shifted solution satisfy the earlier horizon KKT

conditions. Subtracting the stationarity condition (5) of one case from the other and multiplying by ΔP_i produces:

$$\begin{aligned} 0 &= [\nabla f_i(P_i^*) - \nabla f_i(P_i)]^\top \Delta P_i \\ &+ \sum_{j=1}^{N_i} [\mu_{i,j}^* \nabla g_{i,j}(P_i^*) - \mu_{i,j} \nabla g_{i,j}(P_i)]^\top \Delta P_i \\ &+ \Delta \lambda^\top \Delta P_i \end{aligned} \quad (23)$$

The inequality $f(y) - f(x) \geq \nabla f(x)^\top (y - x)$ holds for any convex function f . Applying this rule and using the non-negativity of the μ multipliers:

$$\begin{aligned} 0 &\geq f_i(P_i^*) - f_i(P_i) + f_i(P_i) - f_i(P_i^*) \\ &+ \sum_{j=1}^{N_i} [\mu_{i,j}^* g_{i,j}(P_i^*) - \mu_{i,j}^* g_{i,j}(P_i) \\ &\quad + \mu_{i,j} g_{i,j}(P_i) - \mu_{i,j} g_{i,j}(P_i^*)] \\ &+ \Delta \lambda^\top \Delta P_i \end{aligned} \quad (24)$$

Using the complementary slackness equations (7) this simplifies to:

$$\begin{aligned} 0 &\geq - \sum_{j=1}^{N_i} [\mu_{i,j}^* g_{i,j}(P_i) + \mu_{i,j} g_{i,j}(P_i^*)] \\ &+ \Delta \lambda^\top \Delta P_i \end{aligned} \quad (25)$$

The multiplier non-negativity and primal feasibility inequalities (6) require that the sum term is non-negative implying:

$$0 \geq \Delta \lambda^\top \Delta P_i \quad (26)$$

□

We refer to the value of the right hand side of inequality (26) as the RP indicator for an agent. Any convex agent that is consistent between consecutive horizons must satisfy inequality (26); however, inconsistent agents can also potentially satisfy it. Given the powers P_i and P_i^* , and prices λ and λ^* for two consecutive horizons, the RP indicator is used to identify an agent i as inconsistent if $\Delta \lambda^\top \Delta P_i > 0$.

B. Anticipated Cost Indicator

A second indicator called the anticipated cost (AC) indicator measures how agents anticipate their future monetary costs. This indicator uses the insight, as discussed in the introduction, that agents can often manipulate the receding horizon mechanism by pretending to have high electricity requirements in a future time step: they appear to anticipate higher costs in earlier horizons for particular time steps.

For a given agent i , horizon h and time step $t > h$, the AC indicator $\alpha_{h,i,t}$ is defined as the difference between the anticipated and final monetary costs:

$$\alpha_{h,i,t} := \lambda_{h,t} P_{h,i,t} - \lambda_{t,t} P_{t,i,t} \quad (27)$$

When this value is positive the agent has overestimated what their costs will be, and when it is negative they have underestimated them. These values can be grouped in different ways to further tease out patterns. If the market is a simple linear system we would expect the AC indicator would tend to average to zero over time for an uncertain truthful agent with Gaussian uncertainty.

VIII. INDICATOR EXPERIMENTS

In this section we test the RP and AC indicators on the greedy agent experiments from section VI, and new problem instances where the agent is truthful and uncertain instead of being greedy.

A. Uncertain Agent Problem Instances

Truthful uncertain agents replace the greedy agents in the problem instances. They end up with the same fixed load requirements as the greedy agents, but they do not know this value up front and they have a different estimate for each horizon. We mimic this uncertainty by randomising the estimated load for each time step $t > h$ in horizon h as so: $\hat{P}_t \sim \dot{P}_t |\mathcal{N}(1, 0.5^2)|$, where \dot{P}_t is the actual load that the agent will consume at time t . A final scaling is applied to these powers to make sure that the total power requirements for the whole horizon remains constant. This represents an agent that knows how much it needs to consume, but that is uncertain about when it will consume it. For example, this could represent a house that has a good estimate of its total daily requirements, but who is unsure about when this energy use will take place because of uncertainty about occupant behaviour and solar production. A normal distribution is a reasonable first approximation of the combination of many small uncertain factors that go into making up the total house consumption at each point in time.

Figure 7 shows the solution for an uncertain agent in one of the wholesale market instances. The lines get gradually more opaque with an increase in horizon number. This figure shows that the battery changes its power considerably in response to the variability of the uncertain fixed load.




Figure 7: Horizon solutions for an uncertain agent instance in the wholesale market setting (uncertain agent with 30% of total load, and battery maximum power at 30% of total load).

B. Inconsistency Results

The first experiments apply the RP indicator to the greedy and uncertain agent results. Figures 8 and 9 show how the RP indicator means for 20 instances compare for the greedy (solid lines) and uncertain (dashed lines) agents, over the different market parameter settings. For the wholesale market there is

a clear separation between the greedy and uncertain agent behaviour. Later on we will leverage this in our tests.




Figure 8: Revealed preference indicator for greedy (solid lines) and truthful uncertain (dashed lines) agent in wholesale market setting.




Figure 10: Anticipated cost indicator for greedy (solid lines) and truthful uncertain (dashed lines) agent in wholesale market setting.




Figure 9: Revealed preference indicator for greedy (solid lines) and truthful uncertain (dashed lines) agent in distribution market setting.




Figure 11: Anticipated cost indicator for greedy (solid lines) and truthful uncertain (dashed lines) agent in distribution market setting.

The distribution market setting has more mixed results. If we look back at the cost reductions in Figure 5, we see that the 2% and 5% battery instances are the ones where the greedy agent has most impact. In the 2% battery, 30% greedy share instance the RP indicator is actually less than the uncertain agent's value. The RP indicator might fail to identify manipulation in this situation, but fortunately we find the AC indicator to perform better.

C. Anticipated Cost Results

Anticipated costs are only calculated for an agent in horizon h if it has been identified that it was inconsistent between horizons h and $h + 1$. The results for the greedy and uncertain agents in the wholesale and distribution settings are shown in Figures 10 and 11 respectively. The uncertain agent AC indicator mean remains closer to zero, which is a sign that on average they underestimate costs just as often as they overestimate them.

The AC indicator gives a better separation between the greedy and uncertain agent behaviour than the RP indicator

in these experiments. The 2% and 5% battery results from the distribution setting are now clearly apart.

D. Identifying Agents

The results from the previous section show differences in the behaviour of the greedy and uncertain agents that can be exploited to identify manipulation. As discussed in section VII, this can be done by developing tests that work with various indicators. These tests can work directly with indicator values or be applied to statistical information (e.g., means and variances) collected over time.

It would be trivial to apply such a tests to the 20 day (instance) means we see for the indicators in Figures 8-11. To make it more interesting, assume we have to decide whether the agent is manipulating in each of these 20 days individually. We develop a test and then run it over the 20 days to see how often we identified the greedy agent as manipulative. We design the test so that we get no false positives of the uncertain agents because we are treating their behaviour as being within the acceptable range for an uncertain agent in the market.

The test consists of two very simple linear classifiers that use the agent share of the total load (S as a percentage)

as a variable. The lines in these classifiers were fit to the extreme points of the indicator values produced by the truthful uncertain agent simulations. Agent behaviour that produces a value above one of these lines is classified as manipulative. We focus on the instances where the greedy agent has the biggest impact: 30% and 5% batteries in the wholesale and distribution settings respectively. In the wholesale setting an agent is identified as manipulative when it has a RP indicator over 7519 S , or an AC indicator over 425 S . For the distribution setting these values are 20.0 S and 1.38 S respectively.

For the wholesale and distribution settings, the greedy agent was identified as manipulating in 20/20 days and 19/20 days respectively. More sophisticated tests could be developed to improve on these already very promising results. A great advantage of having this non-intrusive way of testing for manipulation is that it can always be modified or new indicators can be added without having to change the operation of the underlying market mechanism.

E. Scaling to a Real Setting

The indicators we have developed can be used by a receding horizon market operator to reduce the strategies available to manipulative agents. They scale linearly with the number of agents, and only rely on publicly available information. As we have done, the market operator can run simulations with agents that have uncertainty they deem to be reasonable, in order to find an acceptable range of values for the indicators. These simulations could be scaled to much larger systems than we have demonstrated here. We were limited by the computational complexity of calculating optimal greedy agent strategies, which we used to demonstrate the indicators working, but this is not necessarily something that the market operator needs.

IX. RELATED WORK

Our method for calculating greedy agent strategies is related to other work that has looked at equilibria in electricity markets. Hu and Ralph [18] study equilibria in electricity markets with locational marginal prices where each agent solves a bilevel problem to obtain their strategy, and find sufficient conditions for the existence of pure Nash equilibria. Weber and Overbye [19] similarly develop a method for finding Nash equilibria for producers and consumers that have linear price functions, and Li and Shahidehpour [20] consider the case where agents have incomplete information about other agents.

Closer to our approach, instead of searching for Nash equilibria Kozanidis et al. [21] develop optimal bidding strategies for a strategic producer in a single time period market with no network constraints. Compared to these works the problem we investigate is more complicated in certain areas but simpler in others. Instead of a single time horizon, we consider strategies over multiple overlapping horizons given by the receding horizon structure. We also focus more on a prosumer oriented setting where each agent can have a diverse set of preferences and constraints instead of a market dominated by large generator units. As a simplification, we have so far chosen to ignore network constraints.

Mechanism design and game theory have been used in demand response [22]–[24], as well as other network problems including electricity markets and storage adoption [25], [26]. VCG mechanisms have been utilised to achieve incentive compatibility [27], [28]; however, VCG quickly becomes intractable for realistically sized problems, and requires agents to fully disclose their preferences. Tanaka et al. [28] acknowledges these problems and proposes that future work looks at the development of approximate methods. We instead adopt a mechanism which is not incentive compatible (but can be efficiently computed and is budget balanced), and tackle the problem of reducing the practical impact of manipulation.

Mhanna et al. [29] use a scoring rule to charge consumer agents based on both their actual consumption and their deviation from day-ahead allocations. By requiring agents to provide information on their uncertainty, they can reduce the incentive for agents to lie about their requirements over the day ahead. They find it to be “asymptotically” incentive compatible as system size or reported precision increases; however, it does not enable agents to share information they gain throughout the day and re-optimize their allocation in an online manner.

Chapman and Verbic [17] is the only other work we are aware of that has considered the impact of agents manipulating a receding horizon mechanism. Using a clock auction to allocate loads, they discount prices between horizons to give some flexibility for uncertain agents. This discount factor provides a trade-off between allowing agents to recover from uncertain events and preventing manipulation.

Our approach avoids this trade-off in cases where an indicator test can be developed characterise acceptable uncertain behaviour. We found that this is possible with the revealed preference and anticipated cost error indicators, at least when agents have simple forms of uncertainty. These indicators can also be applied to more general settings, such as those where agents are both producers and consumers. To the best of our knowledge, our work is the first to formalise and investigate the significance of manipulation in receding horizon power balancing, and to provide a practical solution for identifying the occurrence of such manipulation.

X. CONCLUSION AND FUTURE WORK

We formally introduced the notion of receding horizon manipulation in a receding horizon market setting, and by developing a strategic agent we empirically identified how much advantage an agent can gain in two representative market settings. We developed indicators for identifying inconsistency and manipulation, which can be used to monitor the interactions of agents in a non-invasive manner. We successfully used these indicators to distinguish agents strategically manipulating the receding horizon mechanism from those that are truthful but uncertain.

Future work will expand these results to a locational marginal price market, i.e. one that takes into consideration network power flows and constraints. For the indicators this should be a relatively straight-forward process of adding additional variables to represent reactive power, while the rest of the theory should remain the same.

Experiments with more diverse agents and diverse forms of uncertainty will be valuable, and with them the development of additional indicators if those presented here are found to be insufficient. The bounds on how much agents have to gain from manipulation could be further narrowed down by considering incomplete information scenarios.

Our experimental work is limited to only considering multiple strategic agents in the cases where they work together as a cohesive coalition. In future work we hope to investigate situations where multiple strategic agents are competing with each other, by searching for (Bayesian Nash) equilibria. In typical settings we would expect this extra competition to reduce the individual rewards that strategic agents can hope to achieve, but it could have a more significant impact on the social welfare than what we observed in our experiments.

Finally, it would be worth investigating if our indicators could still be useful in receding horizon market settings that do not assume convex preferences or utilise marginal pricing.

ACKNOWLEDGEMENTS

Thanks to Archie Chapman for several valuable discussions on game theory and mechanism design.

REFERENCES

- [1] AEMO, “National Electricity Market (NEM),” 2016. [Online]. Available: <http://www.aemo.com.au>
- [2] M. Morari and J. H. Lee, “Model predictive control: past, present and future,” *Computers & Chemical Engineering*, vol. 23, no. 45, pp. 667–682, 1999. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0098135498003019>
- [3] M. Kraning, E. Chu, J. Lavaei, and S. Boyd, “Dynamic network energy management via proximal message passing,” *Foundations and Trends in Optimization*, vol. 1, no. 2, 2014.
- [4] P. Scott and S. Thiébaux, “Distributed Multi-Period optimal power flow for demand response in microgrids,” in *ACM e-Energy*, Bangalore, India, jul 2015.
- [5] B. Kim and R. Baldick, “A comparison of distributed optimal power flow algorithms,” *Power Systems, IEEE Transactions on*, vol. 15, no. 2, pp. 599–604, 2000.
- [6] T. Erseghe, “Distributed optimal power flow using admm,” *Power Systems, IEEE Transactions on*, vol. 29, no. 5, pp. 2370–2380, Sept 2014.
- [7] Q. Peng and S. H. Low, “Distributed Algorithm for Optimal Power Flow on a Radial Network,” *ArXiv e-prints*, Apr. 2014.
- [8] H. W. Kuhn and A. W. Tucker, “Nonlinear programming,” in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley, Calif.: University of California Press, 1951, pp. 481–492. [Online]. Available: <http://projecteuclid.org/euclid.bsm/1200500249>
- [9] W. Karush, “Minima of functions of several variables with inequalities as side conditions,” Masters Thesis, Department of Mathematics, University of Chicago, 1939.
- [10] D. H. Martin, “The essence of invexity,” *Journal of Optimization Theory and Applications*, vol. 47, no. 1, pp. 65–76, 1985. [Online]. Available: <http://dx.doi.org/10.1007/BF00941316>
- [11] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, Jan 2011. [Online]. Available: <http://dx.doi.org/10.1561/22000000016>
- [12] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press, 2007.
- [13] A. Migdalas, P. Pardalos, and P. Värbrand, *Multilevel Optimization: Algorithms and Applications*, ser. Nonconvex Optimization and Its Applications. Springer, 1998. [Online]. Available: <https://books.google.com.au/books?id=zNAwFXUOdgsC>
- [14] P. Scott, “Distributed coordination and optimisation of network-aware electricity prosumers,” Ph.D. dissertation, College of Engineering and Computer Science, ANU, 2016. [Online]. Available: <http://hdl.handle.net/1885/110027>
- [15] L. M. Ausubel, P. Cramton, and P. Milgrom, “The clock-proxy auction: A practical combinatorial auction design,” in *Combinatorial Auctions*. MIT Press, 2006, ch. 5, pp. 115–138.
- [16] Y. Shoham and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. New York: Cambridge University Press, 2009.
- [17] A. C. Chapman and G. Verbic, “An Iterative On-Line Mechanism for Demand-Side Aggregation,” *ArXiv e-prints*, Jun. 2015.
- [18] X. Hu and D. Ralph, “Using epecs to model bilevel games in restructured electricity markets with locational prices,” Faculty of Economics, University of Cambridge, Cambridge Working Papers in Economics, 2006. [Online]. Available: <http://EconPapers.repec.org/RePEc:cam:camdae:0619>
- [19] J. Weber and T. Overbye, “A two-level optimization problem for analysis of market bidding strategies,” in *Power Engineering Society Summer Meeting, 1999. IEEE*, vol. 2, 1999, pp. 682–687 vol.2.
- [20] T. Li and M. Shahidehpour, “Strategic bidding of transmission-constrained genscos with incomplete information,” *Power Systems, IEEE Transactions on*, vol. 20, no. 1, pp. 437–447, Feb 2005.
- [21] G. Kozanidis, E. Kostarelou, P. Andrianesis, and G. Liberopoulos, “Mixed integer parametric bilevel programming for optimal strategic bidding of energy producers in day-ahead electricity markets with indivisibilities,” *Optimization*, vol. 62, no. 8, pp. 1045–1068, 2013. [Online]. Available: <http://www.tandfonline.com/doi/abs/10.1080/02331934.2013.801473>
- [22] A. Mohsenian-Rad, V. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, “Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid,” *Smart Grid, IEEE Transactions on*, vol. 1, no. 3, pp. 320–331, dec. 2010.
- [23] C. Chen, S. Kishore, and L. V. Snyder, “An innovative rtp-based residential power scheduling scheme for smart grids,” in *ICASSP*, 2011, pp. 5956–5959.
- [24] C. Akasiadis and G. Chalkiadakis, “Agent cooperatives for effective power consumption shifting,” in *AAAI Conference on Artificial Intelligence*, 2013. [Online]. Available: <http://www.aaai.org/ocs/index.php/AAAI/AAAI13/paper/view/6193>
- [25] P. Andrianesis, G. Liberopoulos, G. Kozanidis, and A. D. Papalexopoulos, “Recovery mechanisms in a joint energy/reserve day-ahead electricity market with non-convexities,” in *Energy Market (EEM), 2010 7th International Conference on the European*. IEEE, 2010, pp. 1–6.
- [26] P. Vytelingum, T. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings, “Theoretical and practical foundations of large-scale agent-based micro-storage in the smart grid,” *J. Artif. Intell. Res. (JAIR)*, vol. 42, pp. 765–813, 2011.
- [27] P. Samadi, H. Mohsenian-Rad, R. Schober, and V. Wong, “Advanced demand side management for the future smart grid using mechanism design,” *Smart Grid, IEEE Transactions on*, vol. 3, no. 3, pp. 1170–1180, Sept 2012.
- [28] T. Tanaka, A. Cheng, and C. Langbort, “A dynamic pivot mechanism with application to real time pricing in power systems,” in *American Control Conference (ACC), 2012*, June 2012, pp. 3705–3711.
- [29] S. Mhanna, G. Verbic, and A. Chapman, “A faithful distributed mechanism for demand response aggregation,” *Smart Grid, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2015.
- [30] A. Venkat, I. Hiskens, J. Rawlings, and S. Wright, “Distributed MPC strategies with application to power system automatic generation control,” *Control Systems Technology, IEEE Transactions on*, vol. 16, no. 6, pp. 1192–1206, 2008.
- [31] Z. Chen, L. Wu, and Y. Fu, “Real-time price-based demand response management for residential appliances via stochastic optimization and robust optimization,” *Smart Grid, IEEE Transactions on*, vol. 3, no. 4, pp. 1822–1831, dec. 2012.