4
Proposed methodology

In general, there are three alternative or complementary strategies to face uncertainty: (1) to be robust; (2) to be knowledgeable; and (3) to be flexible (Moczydlower *et al.*, 2012). In Chapter 2, we described these strategies, showing the importance of them to make better decisions under conditions of uncertainty. Chapter 3 presented a literature review, reporting some related works, focusing on showing how they incorporate uncertainty and information to make decisions. Within the context of reservoir development, we can summarize the literature review saying that the methodologies available do not give a quantitative value of flexibility through a realizable flow control strategy to a real situation.

The principal focus of this work is on defining the correct problem formulation, determining an optimal development strategy in light of uncertainty and the acquisition of future information, and on the development of efficient methods to determine approximate solutions to this problem. The work does not seek to develop a novel optimization routine, but instead uses the appropriate existing optimization routines as part of a complete methodology. So, in this chapter, we propose a methodology, based on ADP, that is able to define the flow control strategy over time, considering uncertainty scenarios, and the possibility of incorporating the further information acquired to reduce the uncertainty at the time to make decisions. We seek to be robust, knowledgeable and flexible, without requiring an astronomically large number of evaluations, making this approach practical for expensive optimization problems as well.

We apply the proposed approach to value the flexibility of smart wells, defining the optimal flow control valves, in the presence of significant geological uncertainty scenarios, following a combination of both proactive and reactive strategies, permitting much greater flexibility over the timing of valve adjustment and associated measurement acquisition without a corresponding increase in numerical complexity.

4.1. An approach to value flexibility considering uncertainty and future information

To examine the value of flexibility of smart wells we need first to quantify the benefits of this technology, by optimizing the flow control strategy, and then comparing the possible gains against an equivalent case without flexibility (conventional wells). The goal is to control the zones at time intervals (e.g. monthly, quarterly etc.) in order to achieve assigned maximum production for as long as possible. Therefore, in summary the idea of the optimization strategy proposed in this thesis is to seek the optimum control settings of the smart wells, over all time, maximizing the expected NPV under uncertainties and dynamically reacting, adjusting the control settings, with respect of resolution of uncertainties through the acquisition of future information. The need to adjust the control settings arrives from the fact that the acquisition of information reduce uncertainties and consequently alters the ideal control settings. As a result, we have a “playbook” describing the management of valve settings over all time, reacting to the information acquired, and the optimized expected NPV of this flexibility. This playbook also can be used to allow us to value the flexibility in alternative scenarios, without spend optimization.

This approach uses the ideas of ADP (Powell, W., 2010), described in the previous chapter, to reduce the computation burden. Despite this methodology limits the flexibility somewhat so that not all the value of complete dynamic programming is retained – for instance, since the acquired information only affect the decision after it has been made, this policy will not remake the early decision in order to increase the future value. That such losses are often small, and will be more than offset by the increased flexibility that can be feasibly simulated with this approximation.

It begins by optimizing the valve settings over all time, maximizing the expected NPV in the absence of future information. The expectation is made over a set of reservoir models representing the reservoir uncertainty. The valve settings can be adjusted at a discrete set of times. The result is the set of best settings, over all time steps, based only on what is known at time zero, i.e., this is the best proactive strategy.

These settings are then applied to the entire set of reservoir models and future measurements are forecast for the next time step. We then proceed to the next time step, i.e., the next time at which valve adjustments are allowed. At this time we incorporate the information forecast for each reservoir model, potentially reducing uncertainty. The procedure for including future information involves applying cluster analysis to the forecast measurements. The details for the cluster analysis done by the proposed approach are in Appendix A. The notion is that measurements falling within a common cluster are associated with models that are indistinguishable using only those measurements. In other words, the original set of models representing the prior uncertainty is partitioned into smaller sets of models that represent the uncertainty after assimilation of measurement data. This part of the methodology identifies when measurement data are informative. Within each cluster we have reduced uncertainty and should consider a change in valve settings going forward.

For each cluster of models, we determine a new optimal proactive strategy for the future valve settings (past valve settings are not adjusted). This creates a recursion in which an effective, and realizable, strategy can be obtained that keeps the benefits of both proactive and reactive strategies. Since this recursion is performed in the forward direction, the number of required simulations is exponentially reduced compared to the complete dynamic programming solution. Figure 4.1 shows a simplified decision tree for our optimization strategy that considers both model uncertainty and future information. In the red box we provide an illustration of the first step of our optimization.

An illustration of the decision tree for the valuation of control valves with information is shown in Figure 4.2. In this case with three decision points and two measurement points, corresponding to choosing the initial valve settings at time t_0, and then possibly changing the valve settings at two future times, t_1, t_2. Measurements are also taken at times t_1, t_2, with the future valve settings chosen in light of this new information.

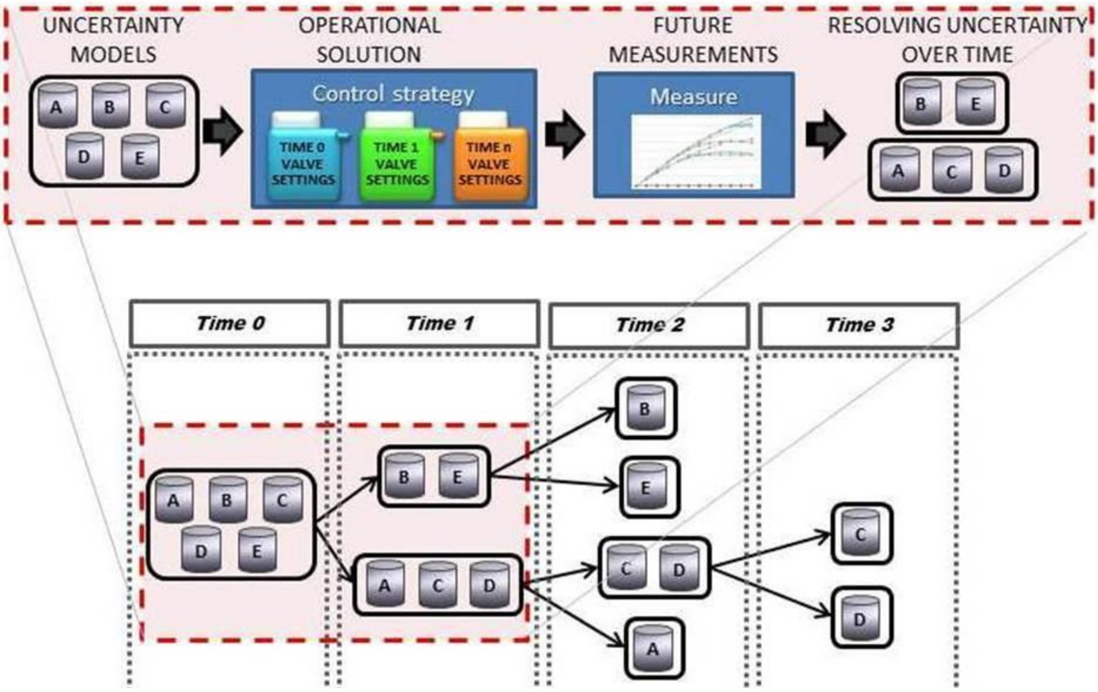

Figure 4.1 Decision tree of our optimization strategy that considers both model uncertainty and future information. The red box is an illustration of the first time step of our optimization.

We implemented this procedure, following the optimization routine proposed by Yeten *et al.* (2002), such that the performance of the reservoir for a particular set of valve settings can be determined via forward simulations. This is accomplished by dividing the entire simulation period into n optimization steps (these steps are distinct from the simulator time steps). The valve settings for the first period (time 0 to time 1) are then optimized. This optimization is performed such that the settings for this period will be the optimum for the entire simulation. We note that this strategy can be applied using different optimization algorithms, when we seek the valve settings that maximize the objective function.

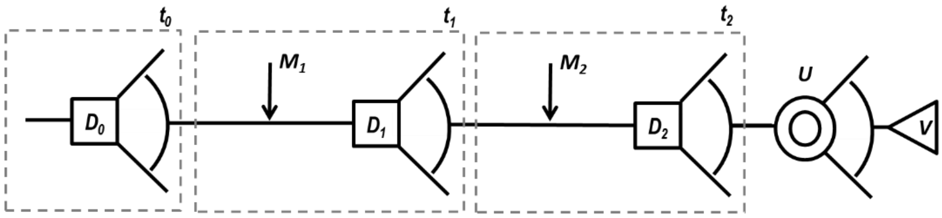

Figure 4.2: Decision tree for a flexible solution with measurement information.

The proposed approach in this thesis is ambivalent about the optimizer used, i.e., we can apply the methodology proposed to value flexibility using any optimizer method to find the optimum controls. Therefore, we can overcome the disadvantage situation about the number of evaluations choosing to use an optimizer that require a reduced number of evaluations to find the optimum controls.

The procedure of re-optimize the settings when information is acquired can lead to repeated evaluations. Therefore, we cannot guarantee that a valve setting will not require evaluation at more than one time, since it may not be an optimum setting at the first time period but could be the optimum setting at the next time period. To solve this problem and consequently reduce the number of evaluation done we implemented a simulation dictionary that saved evaluated solutions. The proposed approach still requiring a large number of evaluations, but most of the expensive reservoir simulations can be simulated independently. We therefore implement all processes using Parallel Programming (Pacheco, 1997). More details about the parallel implementation and simulation dictionary are presented in the Appendix A.

So that we would not require nested optimizations and consequently provide a sufficient reduction in the number of simulations required, we investigated several valuation policies that allow us to effectively use our optimization schemes. The valve settings can be optimized considering different levels of optimization flexibility, varying how often each valve are set during the time horizon. The next Sections describe their variations.

4.2. Timing for optimal control

The most time-consuming part of the computational effort associated with this methodology lies in the reservoir simulations required to determine optimal proactive control strategies. Although this proposed approach is ambivalent to the optimizer method, we choose to use an optimizer that could be efficient for the case of expensive to evaluate functions. The optimization scheme that we employ is based on the Nelder-Mead (downhill simplex) method with the additional use of a radial-basis function (RBF) proxy to enhance the rate of convergence. This approach was described in more detail in chapter 3, it was proposed by Rashid *et al.*. (2013) and provided to us as an optimization library. As is usually the case with optimization schemes, the number of function evaluations (hence reservoir simulations) required for convergence increases as the number of control variables increases. There are a variety of policies that follow this general approach, differing by the degree of flexibility that is afforded to the strategy at each step, i.e., differing by the number of possibility to valve settings adjustment. We describe them in detail in the following sections.

4.2.1. Rolling-flexible

A straightforward strategy allows the valves to be adjusted at each time step during optimization. We refer to this strategy as 'rolling-flexible'. We apply a policy that at each step of the optimization process we seek the optimal flexible valve settings. So at each step we now seek the optimal non-learning flexible policy rather than the optimal static or fixed-valve policy. The optimal valve settings are then given by the follow eq. 4.1-3, with the final valuation given by eq. 4.4.

$$d_0^{\mathrm{RF}}(), *, * = \underset{d_0, d_1, d_2 \in D_{0-2}}{\arg\max} \underset{u \in U}{\mathrm{E}} \left(V | D_0 = d_0, D_1 = d_1, D_2 = d_2 \right), \tag{4.1}$$

$$d_1^{\mathrm{RF}}(d_0^{\mathrm{RF}}, m_1), * = \underset{d_1, d_2 \in D_{1-2}}{\arg\max} \underset{u \in U}{\mathrm{E}} \left(V | D_0 = d_0^{\mathrm{RF}}, D_1 = d_1, D_2 = d_2; m_1 \right), \tag{4.2}$$

$$d_2^{\mathrm{RF}}(d_0^{\mathrm{RF}}, d_1^{\mathrm{RF}}, m_1, m_2) = \underset{d_2 \in D_2}{\arg\max} \underset{u \in U}{\mathrm{E}} \left(V | D_0 = d_0^{\mathrm{RF}}, D_1 = d_1^{\mathrm{RF}}, D_2 = d_2; m_1, m_2 \right), \tag{4.3}$$

$$V^{\mathrm{RF}}() = \underset{U, M_1, M_2}{\mathrm{E}} \Big(V | D_0 = d_0^{\mathrm{RF}}(), D_1 = d_1^{\mathrm{RF}}(d_0^{\mathrm{RF}}, m_1), \\ D_2 = d_2^{\mathrm{RF}}(d_0^{\mathrm{RF}}, d_1^{\mathrm{RF}}(d_0^{\mathrm{RF}}, m_1), m_1, m_2); m_1, m_2 \Big). \tag{4.4}$$

The rolling-flexible policy provides us to plan ahead with our choice of valve settings; in particular the optimal non-learning policy can be obtained from rolling-flexible which is not possible for the rolling-static policy. Even the rolling-flexible policy limits our flexibility somewhat so that not all value is retained – for instance, since the measurements only affect the optimization *after* the previous valve settings have been fixed, this policy will not select valve settings in order to increase the future value from measurements. We suspect that such losses are often small, and will be more than offset by the increased flexibility that can be feasibly simulated with these approximations.

This strategy is the most adaptable, but results in a large number of control variables as it grows linearly with the number of time steps. This, in turn, greatly increases the number of simulations to reach convergence.

4.2.2. Rolling-static

An alternative approach allows the valves to be adjusted just once during the optimization step, thus rendering them constant from that point in time onwards. This approximation was proposed in Yeten *et al.*. (2002).This approximation greatly reduces the number of control variables involved with faster optimizations. We refer to this strategy as 'rolling-static'. Because of the 'rolling' nature of this strategy, it still allows for flexibility. Note that this strategy is still proactive in that it optimizes based on reservoir forecasts, but is also reactive in that it accounts for future information.

When applying the rolling-static valuation policy, we begin with the optimum settings for fixed flow control valves, and use this to fix the initial setting of the valves. With the initial valve setting fixed, we can then determine the range of values for the measurement information up to the first time that the valve settings may be adjusted, and apply the clustering. For each cluster we find the new optimal fixed valve setting for the remainder of the simulation, using this value to fix the second valve setting. This process then continues until all the valve settings have been determined.

Some decision trees that represent the steps of the rolling-static policy for the 3-step problem are shown in Figure **4.3**, where we write $D^{S}_{0\text{-}2}$ for the decision space for all possible *fixed* valve settings at times t_0, t_1, t_2, and similarly $D^{S}_{1\text{-}2}$ is the decision space for all possible *fixed* valve settings at times t_1, t_2. Under the rolling static policy the initial valve setting is given by equation 4.5.

$$d_0^{\mathrm{RS}}() = \underset{d_0 \in D_{0-2}^{\mathrm{S}}}{\arg\max} \underset{u \in U}{\mathrm{E}} \left(V | D_0 = d_0, D_1 = d_0, D_2 = d_0 \right). \tag{4.5}$$

The optimal setting for the valves at time t_1, after clustering the models based on the values of the first measurement is then according to

$$d_1^{\mathrm{RS}}(d_0^{\mathrm{RS}}, m_1) = \underset{d_1 \in D_{1-2}^{\mathrm{S}}}{\arg\max} \underset{u \in U}{\mathrm{E}} \left(V | D_0 = d_0^{\mathrm{RS}}, D_1 = d_1, D_2 = d_1; m_1 \right), \tag{4.6}$$

with the final valve settings given by

$$d_2^{\mathrm{RS}}(d_0^{\mathrm{RS}}, d_1^{\mathrm{RS}}, m_1, m_2) = \underset{d_2 \in D_2}{\arg\max} \underset{u \in U}{\mathrm{E}} \left(V | D_0 = d_0^{\mathrm{RS}}, D_1 = d_1^{\mathrm{RS}}, D_2 = d_2; m_1, m_2 \right). \tag{4.7}$$

With all of the valve settings determined the final valuation generated by following the rolling static policy is according to

$$\begin{aligned} V^{\mathrm{RS}}() = \underset{U, M_1, M_2}{\mathrm{E}} \Big(& V | D_0 = d_0^{\mathrm{RS}}(), D_1 = d_1^{\mathrm{RS}}(d_0^{\mathrm{RS}}, m_1), \\ & D_2 = d_2^{\mathrm{RS}}(d_0^{\mathrm{RS}}, d_1^{\mathrm{RS}}(d_0^{\mathrm{RS}}, m_1), m_1, m_2); m_1, m_2 \Big). \end{aligned} \tag{4.8}$$

The rolling-static policy extends naturally should there be more than three times at which the valve settings may be adjusted.

Although at each step the rolling-static policy requires us to optimize a fixed valve setting, the resulting optimal valve settings will not be fixed. The rolling-static policy allows us to react to future measurement information to maximize the value, and since at each step we always choose the optimal static policy, we automatically have a solution that is somewhat robust to the possibility of valve failure. However, the rolling-static policy does not allow us to plan ahead with our choice of valve adjustments to maximize value and this may lead to an undervaluation. The rolling-flexible policy seeks to address these concerns.

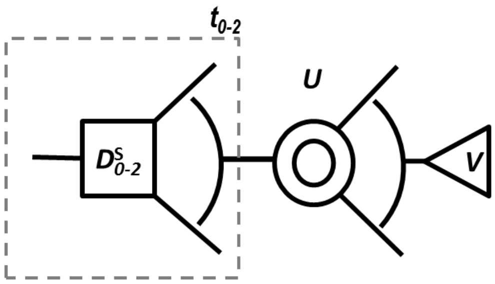

a) Determine optimal fixed valve settings

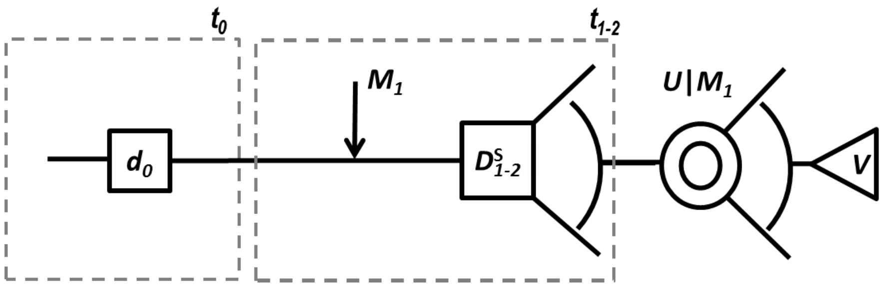

b) Fix initial valve setting, determine measurement values M_1, and determine optimal fixed valve settings for D_1, D_2.

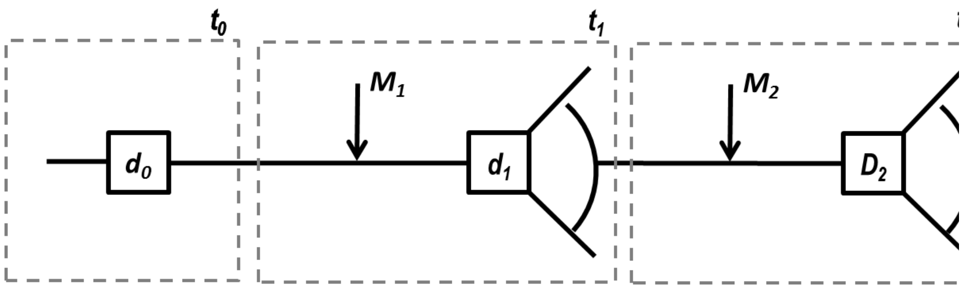

c) Fix d_1, determine measurement values M_2, and find optimal valve settings for D_2.

Figure 4.3: Decision trees used in sequential valuation with rolling static

4.2.3. Rolling-flexible-*k*

A compromise between the maximal flexibility (but high computational expense) of the rolling-flexible strategy and the reduced number of model simulations required by the rolling-static strategy is the so-called ‘rolling-flexible-*k*’ strategy. Here, the valve settings can be adjusted at a limited number of times (*k*), with the valve settings being held constant between these times. While *k* indicates the number of unique valve settings over time, we also need to define how to partition the time horizon accordingly. We propose three possible methods:

• Uniform consecutive: the partitions are the first k-1 time steps followed by a partition comprising all of the remaining time steps;

• Uniform non-consecutive: divides the time horizon into *k* approximately equal partitions;

• Geometric: divides the time horizon into *k* partitions where the size of successive partitions increases approximately geometrically.

Figure 4.4 illustrates these three approaches for *k*=3. Uniform consecutive and geometric partitioning both divide the time horizon so that at the beginning, when the uncertainties are high, the partition sizes are smaller than in later times, when uncertainties are lower. In particular, based on some preliminary tests on simple models, we believe that the geometric partitioning may often be the most appropriate.

Figure 4.5 illustrates a comparison between the rolling-static, rolling-flexible and rolling-flexible-*k* (with *k*=2) strategies. Using the rolling-flexible-*k* (uniform consecutive), the time horizon is divided into 2 parts, and the valve settings can change only twice over the entire time horizon. We can easily observe that using the rolling-flexible-*k* strategy provides increased flexibility compared with the rolling-static strategy and a reduced number of control variables when compared with the rolling-flexible strategy.

Note that all of these strategies can be considered as special cases of the rolling-flexible-*k* strategy. The rolling-static strategy is equivalent to rolling-flexible-*k* strategy with *k*=1. The rolling-flexible strategy is equivalent to rolling-flexible-*k* strategy with *k* equal to the number of time steps between the optimization time and the end of the time horizon.

4.3.
Uncertainty reduction by measurements

As described previously, the proposed approach uses the information derived from measurements to reduce the reservoir uncertainty over time. For that, we consider the forecast information from reservoir simulation, applying the valve settings to the entire set of geological model for the next time step, then we apply cluster analysis to the forecast measurements. The notion is that measurements falling within a common cluster are associated with models that are indistinguishable using only those measurements.

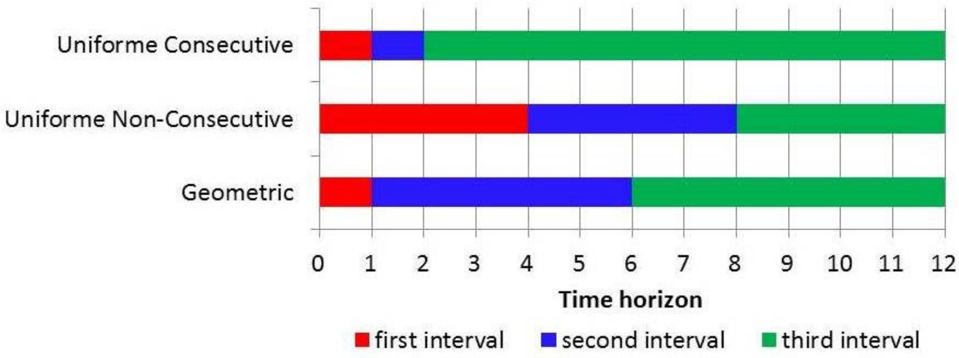

Figure 4.4: Different methods for partitioning the time horizon when using rolling-flexible-*k* (with *k*=3). The different colors represent time intervals where the valve settings are constant.

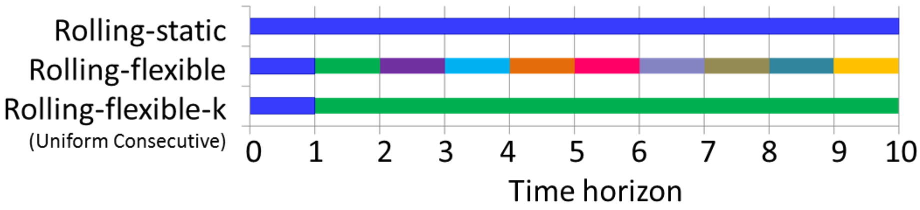

Figure 4.5: Comparison between strategies for the proactive control of valve settings. Each color represents a period of time with constant valve settings

Therefore, we reduce the uncertainty as the original set of models representing the prior uncertainty is partitioned into smaller sets of models that represent the uncertainty after assimilation of measurement data. In other words, once the prior valve settings are known, there will be as many different measurement sets as there are reservoir models describing the uncertainty, thus we apply the rolling methodology with production continuing for infinite time and the clustering is used to group the uncertainty scenarios, reducing the uncertainty at time to make decisions.

Each measurement has an associated “decision resolution”, representing both the accuracy of the measurement and also our policy in how much variation or change is needed in a particular measurement in order to motivate the engineer to make a decision. Measurement accuracy limits our ability to use measurement information to resolve model uncertainty. Even where the resolution of the measurement devices is high, it is unlikely that a decision maker will alter their strategy based on small changes to such measurements. This leads to a coarser granularity of resolution, which we refer to as the “decision resolution”. This denotes the minimum quantum step in a particular measurement needed to make a decision (even if that decision is to do nothing). The need to use a decision resolution that is coarser than the measurement resolution can also be thought of as a consequence of the (by necessity) “limited nature” of the simulation model used to define uncertainty. While our uncertainty models may be sufficient for the purposes of calculating an expected NPV, it might not be sufficient to represent all the possible measurement values with as high a resolution as the measurement itself.

In this approach, we proposed to use the *k*-midranges clustering (Carroll & Chaturvedi, 1998) to partition the uncertainty scenarios. The *k*-midranges is a variation of *k*-means clustering, where the difference in the clustering result originates from the approximation used to find the nearest centroids for the data points, using the midrange as the center of the cluster rather than the mean. Basically, the k-midranges clustering employ the infinity norm (also referred to as the ‘max metric’ and ‘dominance metric’) in finding the range of a cluster, and with this distance metric, the distance of a data point from a centroid is calculated by taking the dominant feature of the difference vector between the data point and the centroid (Malinen, 2015).

In summary, the *k*-midranges algorithm starts by initializing the *k* centroids, and a random selection among the data points is made. Then *k*- midranges consists of two repeatedly executed step: Assign each data point to clusters specified by the “nearest” centroid. Update step: Calculate the new centroids of each cluster of the observations in the new clusters. These steps are repeated until the centroid locations do not change anymore.

To determine the number of clusters required we apply an iterative approach, gradually increasing the number of clusters until the greatest distance between two data points within the same cluster is less than a given decision resolution. As the resulting clustering might still be suboptimal, we apply additional procedures that attempt to reduce the total number of clusters while still maintaining the decision resolution requirement. These procedures consist of testing different starting states for the *k*-midranges algorithm, first through randomly selected initial states and then through attempting to reduce *k* by random removal of previously calculated cluster centers.

Despite to *k*-midranges, as k-means, propose a solution that converge to a local optimum and do not guarantee a global optimality, the *k*-midranges may be best suited to isolating outlying clusters (Carroll & Chaturvedi, 1998). The *k*-midranges produces clusters that tend to be of a similar size, whereas the *k*-means algorithm tends to produce clusters that contain a similar number of data points.

4.3.1. Timing for measurement assimilation

The previous section considered a means to invoke the timing of valve operation for an optimum proactive strategy (considering uncertainty). At the end of each optimization, the ensemble of uncertainty models need to be clustered based on forecast measurements. This is so that the next optimization step can determine the optimal reaction to a future measurement. The number of required optimizations in this approach can be reduced significantly by considering whether it is necessary to optimize the next time step. If there is only one cluster, then uncertainty has not been reduced by the latest forecast measurement and the previously determined optimal settings thus remain valid. The forecast measurements for each successive time step are clustered without re-optimization until two or more clusters are identified, at which time the optimization process is resumed. Thus, the times at which optimizations are conducted are determined by forecast measurements and only when such information results in some uncertainty reduced that is within the resolution of our decision making.

To illustrate this approach, consider the rolling-flexible policy ahead of some future measurement. The valve settings before this future measurement is described by some vector x, and the valve settings, after the measurement has been acquired, are described by some vector y. Before the future measurement, we must optimize the expected NPV (considering all future production), i.e., some function f(x, y), and this gives optimal controls x* and y* for x and y respectively. If the future measurement does not alter the set of models used to calculate the expected NPV then we must optimize the same function f, with the only difference being that the controls before the measurement are now fixed as x*. Since the value of y that optimizes f(x*, y) is still the same y*, the optimal controls do not change.

The aforementioned paragraph is not an exact statement for the rolling static policy, since the controls both before and after the measurement must be equal, and even if, say, $F(x*) = f(x*, x*)$ is the optimum of $F(x) = f(x, x)$ then $f(x*, x*)$ is not guaranteed to be the optimum of $f(x*, x)$. It is also not exact for the rolling-flexible-*k* policies either. Nevertheless, the approximation in these cases should be similar to that which we already make in using these policies in place of rolling-flexible.

With this simple adjustment, the frequency with which we need to re-optimize valve settings becomes intrinsically linked to the arrival of informative measurement data. This suggests the possibility of determining an optimal reactive policy, with an algorithm that imposes no significant constraints on the frequency with which decisions should be made.

Once the prior valve settings are known, there will be as many different measurement sets as there are reservoir models describing the uncertainty. Each measurement will have an associated “decision resolution”, representing both the accuracy of the measurement and also our policy in how much variation or change is needed in a particular measurement in order to motivate the engineer to make a decision. Measurement accuracy limits our ability to use measurement information to resolve model uncertainty. Even where the resolution of the measurement devices is high, it is unlikely that a decision maker will alter their strategy based on small changes to such measurements. This leads to a coarser granularity of resolution, which we refer to as the “decision resolution”. This denotes the minimum quantum step in a particular measurement needed to make a decision (even if that decision is to do nothing). The need to use a decision resolution that is coarser than the measurement resolution can also be thought of as a consequence of the (by necessity) “limited nature” of the simulation model used to define uncertainty. While our uncertainty models may be sufficient for the purposes of calculating an expected NPV, it might not be sufficient to represent all the possible measurement values with as high a resolution as the measurement itself.

4.4. Practical approach

Due to the huge number of uncertainty geological models that can exist, there is a possibility of our optimization model require a number of evaluation that can make this approach intractable. For this reason we can make the optimization procedure less expensive, allowing to optimize the flow control strategy just accounting for representative uncertainty models.

The basic idea about to use representative models is to select a small set of uncertainty scenarios from a large ensemble of uncertainty models, where this selected models can represent the variety of all uncertainty models. The most famous use of representative models is the optimization optimist, realistic and pessimistic where just three uncertainty scenarios are chosen to be optimized together. The next section describe how we can make the proposed approach more practical using a reduced number of evaluations during the valve settings optimization.

4.4.1. Using representative models

Once the option about to use representative models is taken we need to define the number of representative models that will be used during the optimization procedure. This number meaning that each optimization step cannot use more uncertainty scenarios on the evaluation process than the amount previously defined. To choose the representative models that will be used on the optimization process at time 0 we evaluate all uncertainty scenarios using a flow control strategy with the valves fully open (as we don't have a flow control strategy defined yet). So, for example, if the number of representative models was defined equal five, so we use the NPV of each uncertainty scenario to choose the P-10, P-30, P-50, P-70 and P-90 and the optimization for time 0 will be done just evaluating this five models. Once the optimization at the first time step is done, the uncertainty scenarios must be clustered based on the measurement acquisition, so in order to keep the most information possible available we use the optimized flow control strategy at time 0 to simulate and cluster all uncertainty scenarios spare at time 1 (not just the representative models). Then once we split the uncertainty scenarios, each new cluster has their own group of uncertainty scenarios to define the new representative models to be optimized on the current time step. And we follow this procedure for all time steps until the end of the time horizon.

If the number of uncertainty scenarios inside the cluster is less than the number of representative model defined, we can keep the optimization for the remaining time steps evaluating all geological models inside the group.

4.5.
Realizable approach

The key failure of some related works, described on chapter 2, is do not adequately treat reservoir uncertainty, making impossible to apply the optimum control setting in real reservoir development, and consequently making the approaches impossible to be validated. In contrast, the valuation methodology that we are proposing not only determines a final expected value, but also defines via the decision tree an operational plan for future controls (also called by us as playbook). It should even be possible to use the decision tree to control the valves for a new reservoir simulation model that was not used in the original construction of the decision tree, making the proposed approach realizable.

The value of flexibility is calculated by subtracting the expected net present value by the flexible strategy (with smart wells) by the expected net present value by the strategy without flexibility (with just conventional wells). This value represent the economic benefits of the use of smart wells for the reservoir models considered, but how accurate is the final answer?

In order to answer these questions and validate our approach we can split the set of geological scenarios that represent the reservoir model into two groups (drawn from the same distribution), that we call an optimization models group and a test models group. The optimization models group is a collection of geological models, representing the real reservoir uncertainty, used to estimate the value of flexibility and create the playbook (with the flow control settings and reservoir clusters) that maximize the net present value over the time horizon, allowing us to value flexibility. The test models group is another collection of geological models (that was not present on the optimization process), also representing the real reservoir uncertainty and that therefore came from the same distribution as the optimization models group. The decision strategy defined by the playbook (from the optimization process) can be applied to this group of models following the control setting previous optimized according to measurements from the reservoir over time, and obtaining a new valuation. We call this process of using the test models group, a validation test. Such a procedure forms the basis for validation as we present.

Each of the uncertainty scenario from the validation set is simulated up to the first decision node using the initial optimized controls. At each decision node, each validation set then follows the path of the decision tree that corresponds to the cluster whose measurement values are closest to the simulated measurement values for that validation model.

If a particular model has simulated measurement values that differ significantly from all of those represented by the clusters, then this is an early sign that there may be problems with the application of the methodology. In the context of validation we should proceed, with the best course of action possible still being to follow the nearest cluster. We justify the idea of fitting the test models into the nearest cluster because in so doing we can partition the search space, allowing us to always fit a test model into one previous defined cluster. On the other hand, if we have a concept to fit each test model into the previous defined clusters, whose size depends of the measurements resolution, instead of the closest cluster, we could have test models that don’t fit on any cluster, because the holes between them, and consequently we wouldn’t have a flow control strategy to apply. Using this concept of search space partioning we guarantee that always a test model will fit in one of the clusters.

As in the evaluation of machine learning models, the validation and testing can indicate the accuracy of the data. In our approach, if many of the validation uncertainty scenarios show measurements that differ significantly from those represented by the clusters then it suggests that an insufficient number of scenarios were considered during the optimization “training” phase. If we received such anomalous measurement values on applying the strategy to a real field case, we should reexamine the original valuation procedure. Anomalous measurements in this case could be a sign that either insufficient uncertainty scenarios were considered, or that the original model of uncertainty was incorrect.

Our interest in validation stems from concerns about the choice of the number of uncertainty scenarios to be used in training the decision tree (optimizing the flow control strategy) and on the impact of the value of the "decision resolution" that is used to represent the minimum change in measurement value that prompts a "decision" to be taken. The decision resolution must surely be as large as the measurement error resolution, but may be much larger. We are then left with the question of what value to use for the decision resolution, and how to justify such a choice?

With a small value of the decision resolution and a small number of training models we may be able to quickly identify each individual training model, leading to clusters that each contain only a single model and yielding an optimal value, as an "Optimization with clairvoyance". But it is not clear that such an approach would genuinely yield such a value. The value that we obtain could simply be a consequence of an insufficient number of training models leading to `ensemble collapse' at the point of determining an optimal control. Conversely, a large decision resolution may leave a single cluster for the entire development period, yielding a value and strategy identical to that of "Optimization with uncertainty with no use for future information". To address these concerns we need to employ validation.

4.5.1. Uncertainty resolution

As we mentioned, during the validation test it is possible that one or more test models don't have a behavior that had been noted/identified during the measurements assimilation to cluster the models and in this thesis we solve this issue portioning the search space, fitting the test models into the closest cluster. But in order to incorporate more robustness on the clusters, allowing us to reduce the possible number of outliers, we can avoid the uncertainty reduction at some level, i.e., we need to keep such amount of uncertainty in each clustering preventing the full cluster split.

Determining an appropriate value for the decision resolution for every case, reflecting the extent to which the measurements genuinely inform about resolution of model uncertainty and relevance to optimal controls, may be difficult. The difference between the value obtained for the optimization approach and its validation test can be ascribed to ‘ensemble collapse’ in which we are optimizing a small group of uncertainty models, or possibly even a single model, that is not representative of models that yield similar measurement values. Rather than attempt to avoid this ensemble collapse by increasing the number of uncertainty scenarios for optimizing, or by carefully tuning the decision resolution, we can instead proscribe such small groups of models as a step within the clustering algorithm.

We need to keep in mind that increasing the number of uncertainty scenarios considered during the optimization process to define the flow control strategy is possible to get a more robust solution, even with a lower expected value. However in the presence of an expensive evaluation function, as reservoir simulations, increase the number of uncertainty scenarios do not looks good. On the next section we describe how this approach can still being practical even against a large number of uncertainty scenarios to evaluate.

4.6. Reliability by the technical uncertainty

In this work, we want to optimally value and control the smart wells in the presence of uncertainty. Besides geological uncertainty, we also need to plan how to account for technical uncertainty, as equipment failure. As we mentioned on chapter 2, completion failures reduce the field total profitability through decreased revenue and/or increased operational expenditure, consequently when moving into deeper water, the economic penalty for delayed/lost production becomes greater.

If we do not take into account the possibility of such failure when performing a flexible optimization, the resulting policy will have two key shortcomings: 1) it will assign too high a value to the smart completion, and 2) it will not take advantage of the ability of a smart completion to adapt and mitigate when failure occurs. Flexible optimization seeks strategies that are robust in the eventuality of failure by adjusting other valves so that they reduce the consequences of failure.

The reliability is therefore recognized as one of the most important factors for the acceptance of smart well technology. Then, during the optimization of flow control strategy in order to get more robust answer we must account for geological uncertainties and technical uncertainties as well, referring the possibility of equipment failure (equipment reliability).

The exact solution for flexible optimization that includes failure could be provided by including the failure as an additional uncertainty in the same manner that we treated geological uncertainty, i.e., the expected NPV is an expectation over both failure cases and geological uncertainty. The problem with this approach is that the total number of possible failure scenarios can be extremely large, being sometimes impractical. We therefore consider some other approaches that will allow us to examine the impact of valve failure while requiring a feasible number of cases to be evaluated.

The reliable lifetime of many items can be determined statistically from the history of a sufficiently large number of similar items in a population (Veneruso *et al.*, 2000). One way to account for valve reliability is to assign a probability of failure, α, to whenever a valve is adjusted. In our approach, the failure rate of FCVs is defined for a distribution in order to capture the failure behavior at all time horizon. Yeten *et al.* (2004) affirmed that a distribution that is very convenient for representing the life distribution of components where the failure rate varies over time is the Weibull distribution. Nevertheless, our methodology is flexible to use any failure distribution, allowing having results more or less conservatives.

Despite continual improvement, there remains a possibility of valve ‘failure’, either due to failure of the valve itself or the control system. Valve failure leads to loss of future control of the valve and should intervention be required the cost is likely be much higher than the installation cost of the valve. It is important to distinguish between two different modes of valve failure. The most likely mode of valve failure is that the failure occurs before the attempt to alter the valve setting. In this case the valve will be stuck in its previous position. A less likely mode of valve failure is that the valve fails during the attempt to alter the valve setting. In this case the valve does not need to be stuck in its previous position, and for the hydraulically controlled flow control valves it can be stuck in an undesirable position such as fully open or fully closed. It is also harder to mitigate against this mode of failure, since the position at time of failure does not depend on the previous valve settings. The electrically controlled valve removes the potential for this mode of failure as it is not necessary to cycle through the valve positions. In this work we consider the use of electrical valves allowing continuous aperture, and we assume that the failure just can be known when the valve settings are being changed, being stuck at the previous setting.

We might want to resort to physical intervention on the valves so that the losses can be mitigated, paying for the related operation costs, but according Drakeley *et al.* (2003), if the failure consequences are acceptable, the system may be left in the faulty state. On this basis, in this thesis we considered appropriate to classify the equipment as “no repairable” and to measure its reliability performance through survival probability.

Considering that we account for failure every time that the valves must be remotely controlled, we could say that all time steps have some possibility of failure. However, when the valves are installed we assume they have their setting already adjusted so their do not request for remote control. This means that there is no failure possibility at time 0 and for this reason we just account for failure from time 1 to the end.

We divide the approaches to evaluate technical failures in: 1) the approach that considers all valves failing at the first time that they will be adjusted, being completely pessimist and 2) the approach that includes others failures scenarios that can happen with one or more valves in different periods over the time horizon, as follow. In order to consider failure scenarios to examine the impact of failure and/or mitigate losses due the failed valves by the remaining valves we propose six approaches (Figure 4.6): Lucky Optimist, Unlucky Optimist without Learning, Unlucky Optimist with Learning, Lucky Pessimist, Unlucky Pessimist without Learning, Unlucky Pessimist with Learning. They are distinguished by the way that the optimization process are done, including or not possible failures (being pessimist or optimist) and if the failures effectively happen during the life time (being lucky or unlucky).

In this work, the terms “lucky” and “unlucky” are related with there is or not failure during the reservoir development, i.e., “lucky” means that no valves fail during the reservoir lifetime, while “unlucky” means that at least one valve fail at some point. In the other hands, “optimist” and “pessimist” are related with a prior consideration about the possibility of valve failure at the time to define the optimum flow control strategy. Then, “optimist” means that the flow control strategy will be optimized considering that the valves will not fail, while “pessimist” means that the flow control strategy will be optimized considering the possibility of any valve fail and the remaining others valves are able to mitigate the possible losses.

Although the previous explanations about the possibility of failure and its consequences, a first and simplest approach is to ignore it. The Lucky Optimist approach considers an optimistic strategy that do not consider that the valves can fail (because it is optimist) and we don’t evaluate the impact of valves fail because we also assume that we a lucky, and for that reason the fail never occurs. Follow this simple approach is the same as consider a flexible flow control strategy without account for failure in any moment.

On the opposite direction, we have the Unlucky Pessimist with Learning, that means define the flow control strategy accounting for failures and then simulate that this possible failure already happen trying to generate a new strategy to mitigate the possible loses. In that case, if the proactive strategy was right done to consider failures we will find a static flow control strategy, as described on section 4.2.2., that represent the minimum possible gains that we can have with intelligent completion since a fully static strategy do not allows to change the valves for all time steps during the optimization.

The others four remaining approaches are described on the following sections.

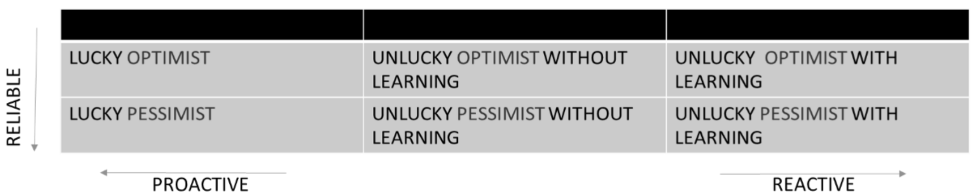

Figure 4.6: Possible approaches that can be incorporate on the proposed approach to incorporate technical failures.

4.6.1. Unlucky optimist without learning

It is possible to quantify the impact of valve failure by examining the expected NPV outcomes for the optimization strategy that does not account for failure. For this, we must to simulate what would have happened if one or more valves had failed during some point of time horizon.

In this approach, we perform a flexible optimization while optimistically assuming that failure cannot occur. We use the Rolling approach, as described in sections 4.2.1-3 with future measurements assimilation to derive the optimal policy for the no-failure case. We then apply this no-failure policy to “unlucky cases” in which failure does occur, and examine the impact of this failure on expected NPV (considering probability of failure with expectation). We name this approach as “unlucky optimist without learning” because we provide a flow control strategy optimistically (without consider fail) and then we evaluate the impact of failure without any reaction on the remaining valves. In summary considering this approach, we are unlucky, because the failure happen, we are optimist because we don’t prepare for failure, and we don’t learn because we don’t alter valve settings after failure.

The failure scenarios can be simulated either in a stochastic process, as Monte Carlo simulation, that allows for various degrees of failure at various times in the life of the well, or in a manual process that allows the decision maker to only analyze failure scenarios that are of interest, for example the worst case where all valves fail. This manual process also allows us to quantify the potential impact of failures and to identify all relevant failure modes and consequences without simulating all failure cases, however it cannot mitigate any possible losses due to failures. An alternative approach is to optimize a policy that accounts for failure, allowing the expected NPV losses due to failed valves to be mitigated by the remaining valves, as will be describe on the next sections.

4.6.2. Unlucky optimist with learning

We recall that to determine the optimal control strategy of flow control valves under uncertainty we frequently have to determine the expected NPV defined as equation 4.9, where c is a vector of the control settings, X is the random variable that describes the uncertainty and V is the resulting NPV determined by reservoir simulation. Using Monte Carlo Simulation (MC) at each iteration a full optimization process is done considering a valve failure schedule and this schedule is defined by the failure probability distribution used.

$$ENPV(X, c) \tag{4.9}$$

In this approach first of all we perform a standard flexible optimization while optimistically assuming that failure cannot occur. And for each iteration we consider a failure scenario, re-optimizing the remaining valves to mitigate possible losses, so even though we be unlucky because the many failure scenarios we have a reactive posture re-optimizing the others valves.

To create a failure schedule we use a quasi-random sampler in order to give for each valve a number that represent the valve fail for all time horizons. The quasi-random (we used Sobol) sampler gives a low-discrepancy sequence of numbers, allowing obtaining a controlled sample. The use of quasi-random sampler require less iteration to sampler rightly a probability distribution, for this reason this sampler allow to accelerate the convergence, reducing the computational cost and improving the efficiency of the methodology (Almeida *et al.*, 2010).

In order to obtain the correct answer using this approach we must to follow the same behavior of a real failure situation, i.e., on real situations we can only change the actual pre-optimized valve settings if one or more valves fail, otherwise we continuing following the pre-optimized strategy. So we use a mask to define the optimum valve settings knowing that a failure will happen on the future, we unconsciously force the optimizer to try mitigate the future losses by subsequence adjustment of the valves. This can make the expected NPV higher and do not return the right answer, because on the real situation they just start to mitigate some failure when this already occur.

In our approach the optimization does not account for failure while this does not happen yet, making the results more consistent with the reality. Thus we define that the first MC iteration is the no failure case and the others left iterations will use the sampler to define the failure schedule. For each iteration we check on the failure schedule when the first failure will occur on the time horizon and we assume that while the failure does not happen the flow control strategy still being the same of the no failure case. Once the one or more valves fail their settings are stuck at the previous setting and all others spare valves will have their settings re optimized in order to mitigate the possible loss. The use of the no failure optimized settings reduces the needless number of evaluations.

So, for example, if the failure schedule for the iteration *i* defines that the first fail will occur on valve *A* at time *t*, all flow control strategy that will be used from time 0 to time *t* are the optimized valves settings by the no failure case. At time *t* we assume that the setting of valve *A* is stuck at the same aperture set at time *t*-1, and all the apertures from the spares valves are re-optimized from time *t* to the end, trying to mitigate the possible losses. Moreover, as the simulation dictionary also described on Appendix B, we create a failure dictionary that save the information about failure schedule and the respective forecast expected NPV. So once a failure schedule is repeated during the MC iteration this case does not need to be re-optimized. This is done for all MC iterations (I_{MC}). Nevertheless, we need to keep in mind that this approach requires a large number of MC iterations, being computationally expensive. The flowchart of this approach is showed in Figure 4.7.

4.6.3. Lucky pessimist

In the cases where the valve failure has a big impact on the evaluation of operation flexibility it is necessary to have an approach that considers the optimization of the valves mitigating the possible losses caused by failure. For this we must to choose the valve settings that maximize the expected NPV accounting for possible failure scenarios in order to mitigate possible losses. So, different of the "Unlucky Optimist", this approach consider the scenarios of failure during the optimization process, not just in the end.

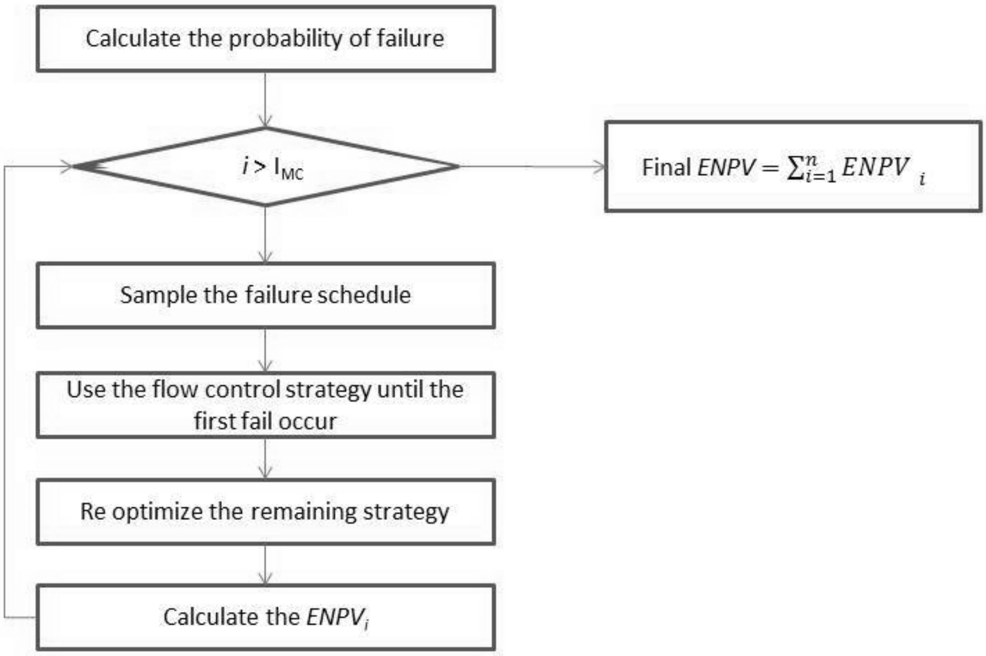

Figure 4.7: Flowchart of optimization process using Unlucky Pessimist.

The "Lucky Pessimist", as we call this approach, adapts the Rolling-flexible approach with no failure by including possible failure in the optimization of the current time step (being a pessimist), but ignoring those potential failures when considering later time steps (being lucky). Then this approach allows the flexible optimization to react in the event of failure, possibly modifying the flow control strategy to mitigate losses in expected NPV due to failure.

This approach is identical to consider the Rolling-static approach for optimization (described on section 4.2.2), because we are looking for the optimum fixed flow control valves, supposing that the no-flexibility for valves adjustment is by some equipment or communication failure. For stuck-in-place failure this is the optimal strategy when failure likelihood is high.

As on the unlucky optimist method, we address the reliability for failure either in a stochastic process, as Monte Carlo simulation, that allows for various degrees of failure at various times in the life of the well, or in a manual process that allows the decision maker to only analyze failure scenarios that are of interest.

In order to try to find a robust solution under geological and technical uncertainties without to ignore the potential failure in later time steps and also react to the failures we propose another approach that use Monte Carlo Simulation to model the failure through the probability distribution, as will be describe on the next section.

4.6.4.
Unlucky pessimist without learning

As we described previously, in this approach we consider the use of electrical valves allowing continuous aperture, and we assuming that the failure just can know when the valve settings is being changed, being stuck at the previous setting.

This approach is similar to “Unlucky Optimist without Learning”, but here we are considering that the valves will be optimized accounting for failures. So we can associate this impossibility to adjust the valve settings due to failure with the inflexibility of the Rolling-Static policy seeking for the optimum static settings to maximize the NPV, since the settings do not change over the time until the acquisition of information reduce the uncertainties.

If at all time steps when the valve are adjusted we seek for the optimum static settings, this valve settings represent the best control at each time considering that all valves can failure later by the lost communication with the valves. So the evaluation of the static valve settings at the end of the first time step represent the evaluation considering that all valves fail at this early time and the valves will not be re-adjusted anymore.

These considerations allow us to know the minimum valuation of the smart control, where the valves aperture are similar to the control of conventional wells, i.e., without flexibility to control the production zones independently and automatic. Moreover this approach does not require extra evaluations, being this evaluation known, because it was included on the optimization process at the first time step. But it does not allow evaluate later failures.