## **SEMIOTIC**LABS

## WHITE PAPER The business case for predictive maintenance

How to calculate the true value a predictive maintenance program will add to your company, with detailed examples and a quick-start guide.



## **Roland van de Kerkhof**

Researcher and lecturer, Tilburg University Founder, Asset Health Dynamics



## Contents

The business case for predictive maintenance	01
--	----

Understanding the value drivers for predictive maintenance	02
The value of time	02
The value of accuracy	03
The value of decisions	05

Methods for calculating the business case	07
Sample method 1: The ROI for an individual asset	80
Example 1	11
Sample method 2: Payback period for a group of similar assets Example 2	13 16
Sample method 3: Upper limit for a group of dissimilar assets	17
Example 3	18

19

A final note	20
	-

## The business case for predictive maintenance

Ramping up predictive maintenance requires serious investments by asset owners and considerable managerial interventions. Having a good sense of the value it will add is therefore one of the preconditions for a successful predictive maintenance program yet calculating the business case isn't easy. In this white paper, I aim to **shed light on the main value drivers** for predictive maintenance and provide **guidance to help you compute the business case** for your own predictive maintenance efforts. In my experience computing business cases, I've typically encountered two groups of people. The first group systematically overestimates the value of predictive maintenance, by assuming long prediction horizons, perfect accuracy, and a wide range of benefits. The second group argues that "we'll never know whether an asset would have failed otherwise"—in which they're generally right—and refrains from calculating the business case altogether. This white paper is for both of them.

The first part of the paper explains three important factors—time, accuracy and decisions—to consider as you think about the value of predictive maintenance. The second part illustrates multiple methods for calculating the business case for predictive maintenance, even under conditions of uncertainty.

# Understanding the value drivers for predictive maintenance

#### The value of time

Predictive maintenance technologies enable an organization to take proactive actions, such as performing targeted maintenance, clustering maintenance activities, and adjusting asset usage. Few of these actions can be performed instantaneously, however (preparing a maintenance activity, for example, takes time), nor can they be initiated at every moment in time. In reality, most organizations have a response time-the time required to respond to a request for action-which depends on the action to be taken, the organizational context, and the timing of the request. The consequence is logical, yet important: the earlier you know an asset is going to fail, the bigger the range of proactive actions you can take.

In fact, for many actions, organizations have both a minimum response time and an optimal response time. Let's take the overhaul of an electric motor as an example. The **minimum response time** is based on an emergency scenario. If you find out right now that the motor is about to break down, how long will it take you to start the maintenance activity? That might require that you stop the production process (wasting product), hire a skilled maintenance technician from a contractor (at a premium), or obtain a spare electric motor (with emergency shipping). If the consequences of breakdown are great enough, it's possible to save money with such an emergency approach. But the motor's overhaul would be much less costly if the organization had more time to react.

The optimal response time denotes how long an organization needs to optimally perform an action. In this example, the optimal response time depends on the time between planned production stops, the scheduling horizon for maintenance technicians, and the standard delivery time for electric motors. If production is stopped once every month, for example, and the scheduling horizon and delivery time are three weeks, the optimal value of preventive overhaul can be derived if the organization knows more than a month in advance that the motor is about to break down. The technology's prediction horizonhow far in advance a prediction system produces a correct prediction—therefore determines the value that can be derived. This idea is visualized in figure 1.

The further you can see into the future, the more proactive responses you have available and the better you can perform each one.

TIP



Figure 1. The value of time: an example of a relationship between a condition monitoring technology's prediction horizon and the potential value you can derive.

#### The value of accuracy

Even if you have all the time in the world, few predictive maintenance technologies are capable of perfectly predicting failures neither from the start nor over time. Most new applications require learning—by machines, by humans, or both—while over time the predictive performance is subject to changes in the asset itself (e.g., modifications) and its operational context (e.g., process or product changes).

Two important performance indicators for a predictive maintenance technology are its **sensitivity** and its **specificity**. The sensitivity, also known as the true positive rate, indicates the percentage of failures that are identified

beforehand (providing the organization sufficient response time). Specificity, or the true negative rate, indicates how well the technology is able to identify that an asset is *not* about to fail. The higher the specificity, the lower the number of false alarms. The higher the sensitivity, the lower the number of unexpected breakdowns. Together, the sensitivity and specificity determine the technology's **accuracy**: the percentage of failures and non-failures that are correctly identified as such.

	Actual condition: positive (failure)	Actual condition: negative (no failure)
Predicted condition: positive (failure)	true positive (TP)	false positive (FP) (type I error)
Predicted condition: negative (no failure)	false negative (FN) (type II error)	true negative (TN)
	SENSITIVITY:	SPECIFICITY:
	TP	TN
	TP + FN	FP + TN

Figure 2. The value of accuracy: classifying predictive sensitivity and specificity.

To calculate the business case for a new predictive maintenance technology, we have to take into account that its accuracy is not perfect. Especially for complex assets with multiple failure modes and degradation mechanisms, business case analyses should incorporate the probability of missing an upcoming failure and the probability of raising a false alarm. In addition, it should be noted that for many assets, the current accuracybefore implementing the new predictive maintenance technology—is rarely zero. Anomalies and upcoming failures can, for example, be detected during visual inspections, functional tests, and via production interference, although the prediction horizon of these methods is typically lower than with predictive maintenance technologies. Sound business cases therefore focus on the difference in accuracy between the old and new situations.

Few predictive maintenance technologies are capable of perfectly predicting failure.

TIP

For most assets, the current prediction accuracy isn't zero, even if you aren't yet using predictive maintenance technologies.

TIP

#### The value of decisions

Almost by definition, predictive maintenance is intended to reduce the cost of maintenance, by enabling you to skip scheduled maintenance activities, prevent unexpected breakdowns, reduce the frequency of inspections, cluster maintenance activities, or perform focused maintenance. This value is generated by **making decisions that are better informed**. But insight into the current and future state of assets can also benefit other stakeholders in the organization, such as the production department—by reducing energy and materials usage, increasing availability, reducing slowdowns and reducing quality losses—and the project department by extending assets' useful life. Table 1 summarizes common value drivers for predictive maintenance, including the range of realized benefits I've observed in practice (in percentages).

Value driver	Observed percentages	Beneficial for
Reduced maintenance costs	-10% - 50%	Maintenance department
Reduced capital expenditure	-10% - 50%	Project department
Reduced safety & environmental risk	0% – 50%	Many stakeholders
Reduced operational costs	0%-10%	Production department
Increased overall equipment effectiveness	0% – 25%	Production department

Table 1. Common value drivers for predictive maintenance, including benefits observed in the field.

Predictive maintenance can also lead to substantial benefits for other stakeholders, beyond the maintenance department.

TIP

Two things should be noted here. First, while each predictive maintenance use case can have multiple value drivers, only one or two are generally dominant. For example, if predictive maintenance is used to extend an asset's useful life, the cost of maintaining the asset (and the associated risks) tend to increase. If predictive maintenance is used primarily to maximize the asset's uptime, maintenance costs tend to remain stable.

## Understanding the value drivers for predictive maintenance

Of course, there are examples in which predictive maintenance results in less unnecessary and time-consuming maintenance, thereby automatically increasing the asset's overall equipment effectiveness (OEE) and limiting the number of risky maintenance activities. In these situations, the asset's context mainly determines which benefit is dominant: for some contexts, uptime is much more valuable; for others, capital expenditure or the cost of maintenance.

Second, in my research I've observed that it can take up to several years before the information provided by the predictive maintenance technology is used in decisionmaking, especially if the technology is not yet perceived as "proven" and the decisions carry risk. This reduces short-term benefits, as the technology only provides value if the organization's decision-making is improved.

During that time, it's possible for the predictive maintenance technology to generate "negative benefits." If the costs of maintenance and capital expenditure haven't yet declined, the initial investment in purchasing and installing the technology and the operational cost of using it to perform measurements and analyses can actually increase overall capital expenditure and maintenance cost. Moreover, if the predictive maintenance technology generates many false alarms, the number of maintenance activities might actually increase, further raising the cost of maintenance.

While predictive maintenance can have multiple benefits, typically one or two value drivers are dominant for each

use case.

TIP

The benefits of predictive maintenance technologies are realized via improved decisionmaking.

06

TIP

## Methods for calculating the business case

There are many ways to calculate a business case, as well as a wide variety of outcome variables. The most dominant outcome variables are the **return on investment** (ROI) —the ratio between the financial gain an investment produces and its cost—and the **payback period**, or time it takes to recover the investment's cost.

Selecting the appropriate method depends, among other things, on the technology's use case—to monitor an individual asset, a group of similar assets, or a group of dissimilar assets—and how important timing is. If it doesn't matter when costs are incurred and revenues are earned, you can simply use averages to calculate the business case (such as the mean time between failure, average cost of breakdown, and so forth). If timing does matter, such as when computing a payback period or an ROI with a discount rate, you'll need to run simulations. In this section, we'll look at three sample methods for calculating the business case for predictive maintenance.

	Individual asset	Group of similar assets	Group of dissimilar assets
Timing <i>does not</i> matter	Example 1: ROI without discount rate		Example 3: Upper limit
Timing <i>does</i> matter		Example 2: Payback period	

Table 2. Three sample methods for calculating the business case.

## EXAMPLE: SETUP

For our three sample calculations, we'll use a fictitious company called Solid Steel. Solid Steel is a multinational steel manufacturer with production plants in 25 countries and annual production of 25 million metric tons.

You're working at the company's Dutch plant, Solid Steel NL, a site that produces steel around the clock, with the exception of a few scheduled maintenance stops each year. The plant has many rotating assets, including 4 compressors, 20 fans, 100 pumps, and 200 conveyors.

Two weeks ago, your management indicated that lately the pumps have been having issues, particularly several of the 55 centrifugal pumps. Of these 55 centrifugal pumps, 35 are already incorporated in a manual vibration monitoring program (which measures them every six weeks). The 15 most critical pumps are manually checked for lubricant once every year.

You've recently heard about an affordable automated monitoring system, so you're setting out to identify the pumps for which the new system would be valuable.

#### Sample method 1: The ROI for an individual asset

Let's start with the basics. The return on investment depends on the difference in costs and gains between the present state and the future state. All features that are unaffected can be left out of the business case—and the simpler the business case, the better. Experience shows that in most business cases, only a small number of variables typically between three and five—have a strong impact on the business case as a whole.

The costs of predictive maintenance can be subdivided into initial costs and recurring costs. Initial costs are incurred to implement the new system, such as costs for engineering, procurement, installation and training. Recurring costs are those incurred by the system's ongoing use: for example, the costs of inspections, analyses, and management. If these activities are outsourced, the recurring costs will be aggregated in a monthly or yearly subscription fee.

In a business case analysis, costs are generally the easiest part to identify, as they are either specified by the vendor of the predictive maintenance technology or service, or can be retrieved from earlier use cases.

Methods for calculating the business case

The main challenge lies in estimating the benefits of predictive maintenance. Here, a lot of uncertainty arises: it's unknown how often an asset will break down, how well the predictive maintenance technology will perform, whether a modification will be implemented that extends the asset's lifetime, and so on. To make it less complex and reduce the bias in estimation, **I recommend the following procedure:** 



**Gather a team** of the asset's maintenance engineers, the predictive maintenance technology specialist, and you, your company's brand-new business case specialist.



**Start with the direct benefit** from preventing failures. The best way to do this is by decomposing the asset's failure into failure modes. If an asset has 10 ways of failing, write them all down, and assess for each failure mode whether the predictive maintenance technology will improve the sensitivity. If so, for that failure mode:

Estimate the **mean time between failures** (MTBF). You can base your estimate on the manufacturer's manual, existing reliability data, expert judgment, or a combination of these.

Estimate the current **sensitivity** and the new sensitivity using the predictive maintenance system. Both the minimum response time and the optimal response time can be used to estimate the sensitivity and corresponding costs, but be consistent.

Estimate the **costs** for two scenarios: (a) the failure was not foreseen, and (b) the failure was foreseen. Include whatever costs are relevant for your case: the costs of maintenance, the opportunity costs of lost production, the costs of environmental damage, and so on.

h

If variables are uncertain and it's hard to specify averages, use ranges.

Let the specialists collectively come up with a minimum ("It's very unlikely the real value is lower than ...") and a maximum ("It's very unlikely the real value is greater than ...").

Later on, test how substantial the effect is on the business case. If it turns out one of these variables does have a big effect and the range is large, it might be wise to search for additional information to reduce the width of the range.

TIP



a

h

Next, discuss **additional value you'll derive** from the new predictive maintenance system. To what extent and how will it affect inspection costs, periodic maintenance costs, operational costs, operational revenues, and so on? Calculate the resulting benefit separately for each source of value. Common sources of value are:

**Reducing the frequency of other inspections**, if the new technology replaces visual inspections partly or altogether.

**Extending periodic maintenance intervals**, if periodic maintenance is skipped when the predictive maintenance system indicates the asset's condition is still okay, or if periodic maintenance is stopped altogether.

**Reducing energy usage**, if the new system helps you identify and solve energy wastage earlier.

**Extending the asset's lifetime**, if degradation is identified and solved earlier, preventing further degradation, or if the new technology generates insight into the sources of degradation, and these can be mitigated.



a

Last but not least, discuss the **cost of alarms** (both true and false). How many alarms do you expect each year? And how will you respond to them? Discuss the process and identify what costs it will incur. It's common for:

alarms to trigger **additional inspection**, either to validate the alarm or to diagnose what maintenance needs to be performed. If these diagnoses would not have been performed otherwise, they count as additional costs.

alarms to trigger **maintenance actions** (otherwise, the value at step 2 is limited), including unnecessary ones—that is, maintenance that's performed too early. After validation (step 4a), what percentage of alarms still lead to unnecessary maintenance?

### EXAMPLE 1

One of the critical centrifugal pumps has recently broken down unexpectedly, stopping production, and has had relatively high maintenance costs over the past couple of years. So you've decided to start your business case analysis here.

Following the procedure above, you've first requested an indication of the costs from the predictive maintenance technology's supplier:  $\leq 1,000$  to install and  $\leq 500$  per year for the monitoring service. Then you created a table to determine the direct benefits from preventing failures. Table 3 shows part of this table, for three common pump failure modes: impeller failure, bearing failure and seal failure. Based on past reliability data and the maintenance engineer's judgment, you expect the MTBF for these failure modes for this pump to be 10, 8 and 5 years, respectively. This pump's unexpected breakdown is capable of shutting down the production line, causing a production loss of approximately  $\leq 25,000$  per hour.

Most upcoming impeller and bearing failures in this pump have been detected early on through your manual vibration monitoring and lubricant analysis, but the new system can identify impeller unbalance several months earlier, enabling earlier preventive maintenance. That reduces consequential damage to the bearings and increases the likelihood of scheduling repairs during an already planned

maintenance stop. The new system is thus more accurate (thanks to more frequent measurement) and enables you to derive more value from proactive responses (thanks to a longer prediction horizon).

	Sensitivity		itivity	Cost-of-failure scenarios			Average
Failure mode	MTBF	Old	New	Not foreseen	Foreseen old	Foreseen new	annual difference
Impeller failure	10 years	80%	93%	€57,000	€7,000	€5,000	€836
Bearing failure	8 years	75%	95%	€27,000	€2,000	€2,000	€625
Seal failure	5 years	10%	10%	€58,000	€8,000	€8,000	€0

Table 3. Calculating the direct benefits from preventing failures, based on the costs of maintenance and downtime.

Then there are the secondary benefits. According to your maintenance engineer, the pump is cavitating about 25% of the time, resulting in suboptimal efficiency and degradation to the impeller blades, the bearing and the seal. If the pump's operators have real-time insight into the cavitation, they can reduce the flow to steer the pump back toward its best efficiency point. By adding this measure, the maintenance engineer expects the pump will cavitate only 10% of the time.

Your energy usage data reveals that the pump uses around 10% more energy when cavitating, so the reduction in cavitation will produce a 1.5% reduction in the pump's annual energy usage. If the pump is powered by a 45 kW electric motor and is operational 95% of the year, this will save on average  $\in$ 1,205 per year. Moreover, decreasing the amount of time the pump cavitates is expected to reduce the frequency of impeller, bearing and seal failures by 33%. This saves on average  $\in$ 3,533 per year for reduced seal failures and an additional  $\in$ 423 for impeller and bearing failures (calculating this requires updating table 3 with the new MTBFs).

Together, the annual benefits in this example are  $\in 836 + \in 625 + \in 1,205 + \in 3,533 + \in 423 = \in 6,622$ . Given the  $\in 500$  annual service fee, the business case is positive in the long run if and only if the additional costs for unnecessary maintenance (in step 4) average less than  $\in 6,122$  per year. Since the new system has a specificity of 92%—meaning it rarely sounds a false alarm—this seems quite feasible.

#### Sample method 2: Payback period for a group of similar assets

If the organization is interested in the payback period of the investment, the timing of events becomes important. Let's assume an asset fails once every 10 years, the new predictive maintenance technology increases the probability of detection by 50%, and the organization is only willing to invest if the payback period is two years or less. With a constant hazard rate,\* the asset has a  $\pm 20\%$ chance of failing in the first or second year, of which an extra 50% is now detected.

If you rely on the averages per year to calculate the payback period, you'll systematically overestimate the likelihood of achieving it. So there are two ways to go: manual probability calculations or computerized simulations. This white paper is too short to write a complete guideline for either of these methods, so here I'd like to share some insights into one of them: simulation.

A **simulation** is an imitation of a process or system—such as an asset's maintenance process—over time. By creating a simulation model for an asset or group of assets, it becomes possible to observe what is likely to happen in the future. The future is uncertain, yet some scenarios are more likely than others. By running a simulation repeatedly (a thousand times, for example), you can generate a probability distribution for each outcome. The question then becomes: of these one thousand simulated futures, how many of them had a payback period shorter than two years?

There are several software packages you can use to simulate a process, such as Arena, AnyLogic, R, and even Microsoft Excel. Each has its own advantages, disadvantages and language. For our example I'll use Sysdea, a system dynamic software package that allows for the quick development and visualization of simulation models. Figure 2 shows a sample Sysdea model for simulating the value a new predictive maintenance technology will contribute to a factory.

Once you've developed a basic structure, it becomes relatively easy to redo the simulation for other predictive maintenance technologies. For each use case, you need to determine the variables involved (most change from one use case to another) and

\* In industrial settings,  $\pm 87\%$  of assets arrive at a constant hazard rate (Plucknette, 2005). Most assets experience more hazards at startup ( $\pm 67\%$ ), followed by assets that have a constant hazard rate from the start ( $\pm 13\%$ ) and assets that have fewer hazards at startup ( $\pm 7\%$ ). Only  $\pm 13\%$  of assets have an increasing hazard rate over time.



![](_page_15_Figure_3.jpeg)

check the relationships between them (most remain the same). In addition, once the basic structure has been developed, it becomes relatively easy to increase the number of assets, especially if the assets and their maintenance processes are similar to each other: they have the same hazard rate, similar consequences of failure, similar costs of maintenance, and so forth.

#### So when are simulations a good way to go?

- If the decision is very important.
- If timing matters.
- If the system is complex (for example, when variables are interrelated).
- If you're interested in getting more detailed insight into the system.
- If the number of similar assets is medium (more items means less chance of bias, the thing you're trying to eliminate through simulation).

The quality of a simulation, as well as the time it takes to construct, depends on the skill of the builder. If you're not proficient in making simulations (yet), I recommend finding someone in your organization who is. I've seen people in business intelligence, data science, reliability engineering and process optimization departments using simulations, so these might be a good place to start your search.

TIP

## EXAMPLE 2

Your manager is glad to hear that the predictive maintenance business case is positive in the long run for the first centrifugal pump. But he's only willing to invest if the payback period is likely to be shorter than two years—a probability of 80% is sufficient. So now you're consulting with a simulation specialist from the company's business intelligence department to compute a distribution for the payback period. The underlying question is: how many of your centrifugal pumps need to use the new technology to reach an 80% chance of recouping the cost within two years?

The costs are easy to compute: with every additional pump, the initial costs increase by  $\in$ 1,000 and the annual costs by  $\in$ 500. The cost per pump is thus  $\in$ 2,000 for the first two years. The benefits are harder to compute, so here you consult the maintenance engineer again. She tells you that the average MTBF of the three failure modes is slightly longer than for the first centrifugal pump–12 years for impeller failure, 10 years for bearing failure and 8 years for seal failure, on average–and the costs of unexpected breakdown are slightly lower on average, since not all the centrifugal pumps are critical. The average reduction in energy consumption is estimated conservatively at  $\in$ 500 per pump per year. Over the course of two years, this results in a cost reduction of  $\in$ 1,000 per pump.

This means the desired payback period of two years will be achieved if enough breakdowns have been prevented to reduce costs by (number of pumps) \* ( $\varepsilon 2,000 - \varepsilon 1,000$ ). The higher the number of pumps incorporated into the new predictive maintenance program, the greater the likelihood that one or more will fail in the first two years.

Let's test it for 10 pumps. With a hazard rate of 0.1 (which corresponds to a MTBF of 10 years, given a constant hazard rate), the probability of having at least one failure in the first two years is ±86% (using the binomial distribution). If preventing one failure is sufficient to cover the predictive maintenance expenses for 10 pumps, the payback period is sufficiently likely (> 80%) to be shorter than two years.

But you also need to take into account that the initial sensitivity wasn't 0%—in fact, it was already quite high for the first pump—so you need to use the difference in sensitivity. If the difference in sensitivity is 50%, you need about twice the number of pumps to prevent a failure you wouldn't have prevented otherwise. If the difference in sensitivity is 25%, you need four times the number of pumps, or 40. In this case, because many of the pumps aren't being monitored at all yet, your team determines that the average difference in sensitivity is 33%. The business case for the group of pumps will thus be positive in more than 80% of the simulation runs if the benefit of preventing a failure is at least 30 \*  $\leq$ 1,000 =  $\leq$ 30,000 (including the operational benefits from increased uptime, and so forth).

Methods for calculating the business case

If the average business case for the first two years isn't positive, increasing the number of assets won't help you reach the 80% threshold. Increasing the number of assets only increases the likelihood of failures taking place in the first two years. Predictive maintenance is just one way an organization can increase OEE, reduce maintenance costs, and so forth, so only a fraction of the full potential for improvement can be realized through predictive maintenance.

TIP

## Sample method 3: Upper limit for a group of dissimilar assets

TIP

While the business case for a group of dissimilar assets can rely on the same methods as the business cases for individual and groups of similar assets, it becomes increasingly time-consuming to collect the data with each new type of asset added.

Fortunately, if the assets are operationally connected, such as in a single unit or plant, it's possible to create a quick approximation of the achievable benefits. The trick here is to first assess the potential to improve the plant's value drivers, such as maintenance costs, capital expenditure, OEE, and so on, and then calculate the impact the predictive maintenance technology is expected to have on these value drivers. When a factory's management looks at adopting a broader predictive maintenance program, a new question becomes relevant: "How much can this plant gain from applying predictive maintenance, and thus how much are we justified in spending on it?" This analysis can be performed for a single predictive maintenance technology or for a group of them.

Typically, organizations like to invest in a limited number of systems that can be applied over a wide variety of assets, since each new system entails startup costs (such as gaining proficiency and building a relationship with the supplier). The more widely each predictive maintenance system can be used—especially when critical assets are involved—the more value the company can capture from predictive maintenance.

## EXAMPLE 3

The new predictive maintenance system has proven its value on the 30 pumps it's monitoring at Solid Steel NL, and the plant manager has asked you to help determine whether it makes sense to scale up to other assets.

Over the past couple of years, the plant's performance has been relatively stable. Its OEE has fluctuated around 55%, maintenance costs average €25 million per year, and the plant spends €50 million per year on energy. Yet according to global steel industry benchmarks, the best performers have an OEE of 80% and maintenance costs are just 2.5% of their ARV (asset replacement value).

You put together an experienced team, including the plant manager, predictive maintenance specialists, and process and reliability engineers, and start by estimating the plant's theoretical optimum. Since it's quite an old factory, you're expecting its optimal OEE might reach 75%, slightly lower than the benchmark's best performers. Compared to the current 55% OEE, this is a 36% improvement.

The factory currently has a turnover of roughly  $\in$ 200 million per year. In the optimal scenario, this could go up to  $\in$ 272 million each year, an increase of  $\in$ 72 million per year. Equipment unavailability is the major reason for the plant's current low OEE (availability = 0.7, productivity = 0.9, quality = 0.9). The team estimates that a full-fledged predictive maintenance program can raise availability to 0.8, but won't have a big impact on productivity and quality. That will raise your OEE to 65%, resulting in annual turnover of  $\in$ 236 million, a  $\in$ 36 million increase per year.

Using a similar analysis, the team estimates that an optimal predictive maintenance program will reduce annual maintenance costs by  $\in$ 5 million (to  $\in$ 20 million per year) and energy usage by  $\in$ 7.5 million (to  $\in$ 42.5 million). The upper limit on the factory's benefit from predictive maintenance is thus  $\in$ 36 +  $\in$ 5 +  $\in$ 7.5 =  $\in$ 48.5 million per year.

In this scenario, Solid Steel NL's management has quite some room to develop its predictive maintenance program!

## **Quick-start guide in four simple steps**

![](_page_20_Figure_2.jpeg)

## A final note

The example in this paper is fictitious; the methods employed are not. I hope the insights I've shared here will spark your creativity and help you assess the business case for predictive maintenance in your own company.

![](_page_21_Picture_3.jpeg)

Roland van de Kerkhof

Passionate about solving the problem of unplanned downtime, Semiotic Labs uses AI-driven electrical waveform analysis to create smart condition monitoring solutions that increase productivity while saving customers' time and money.

www.semioticlabs.com

## Contact

Lars Ligtenberg Sales development representative lars@semioticlabs.com +31 653 325 256

White paper | The business case for predictive maintenance

![](_page_21_Picture_10.jpeg)