Australian Government
Department of Agriculture, Water and the Environment

Discount rates and risk in the economic analysis of agricultural projects

Simon Hone, Peter Gooday, Ahmed Hafi and Jared Greenville

Research by the Australian Bureau of Agricultural and Resource Economics and Sciences

April 2022

Cataloguing data

This publication (and any material sourced from it) should be attributed as: Hone, S, Gooday, P, Hafi, A & Greenville, J 2022, *Discount rates and risk in the economic analysis of agricultural projects*, ABARES Working Paper 22.02, Canberra, April, DOI: https://doi.org/10.25814/pdwp-6y91. CC BY 4.0.

ISSN 189-3128

Australian Bureau of Agricultural and Resource Economics and Sciences (ABARES)
GPO Box 858 Canberra ACT 2601
Telephone 1800 900 090
Web awe.gov.au

Acknowledgements

We would like to thank Tom Kompas and Mark Harvey for their valuable feedback on an earlier version of this report.

Acknowledgement of Country

We acknowledge the Traditional Custodians of Australia and their continuing connection to land and sea, waters, environment and community. We pay our respects to the Traditional Custodians of the lands we live and work on, their culture, and their Elders past and present.

Key points

- There is considerable debate around discount rates and the treatment of risk in economic analysis.
    - Since the 1980s, most government economic appraisal guidelines in Australia have adopted a consistent approach set out in government economic appraisal guidelines.
    - However, support for this position has eroded over time.
    - Critics have argued that the discount rate should be updated to reflect changing economic conditions and suggested, more fundamentally, that the theoretical basis of the standard approach is flawed.
    - If these criticisms hold, we are likely to be underestimating the merits of long-term projects versus short term projects and the merits of projects that reduce risk versus projects that increase risk, with implications for the quality of investment decisions.
    - This report explores both sides of the debate to provide rigorous and practical guidance on how economists should approach discount rates and the treatment of risk in their analysis of agricultural projects.

- We need to discount future costs and benefits.
    - One of the motivations for discounting is to account for the opportunity cost of capital.
    - We do not want to invest in a project if it means forgoing a better alternative.
- The standard approach is to discount based on the (real) expected return on the alternative investment.
    - There are several ways to define the alternative investment and this affects the choice of discount rate.
    - Government economic appraisal guidelines tend to define the alternative as an investment in private assets across the economy.
    - The historical long run return on private investment across the Australian economy has been around 7 per cent (or slightly higher).
    - This is the rationale for the recommended 'central case' discount rate of 7 per cent in most guidelines.
- But the standard approach is not always a good approximation.
    - In particular, it does not account for differences in riskiness between the proposed project and the alternative investment.
    - For example, a zero-risk agricultural project with an expected return of 5 per cent might be preferred to a medium-risk alternative investment with an expected return of 7 per cent.
    - However, the standard approach will always recommend the medium-risk alternative investment, given the higher expected return.
- Some economists argue that a pragmatic alternative to the standard approach is to adjust the discount rate to account for project risk.
    - However, it only provides a good approximation under restrictive conditions.
- A better approach is to address the time value of money (value of a dollar in the future relative to a dollar now) and risk separately.
    - Where the costs and benefits are not known with certainty, estimate the certainty equivalents for the agricultural project in each year.

- Then discount the certainty equivalents based on the (real) risk-free discount rate.

- To demonstrate the theoretically correct approach, we revisited a previous ABARES biosecurity application.
    - Under the theoretically correct approach, citrus canker is estimated to cost growers $320 million in present value terms over the next 50 years.
    - By contrast, under the standard approach, citrus canker is estimated to cost growers just $80 million.
    - Hence, the standard approach underestimates the costs by hundreds of millions of dollars in the case of citrus canker.
    - This can have real world consequences, including failing to make worthwhile investments to limit the arrival or spread of citrus canker because we are substantially underestimating the avoided costs from such investments.
    - The application also shows that the theoretically correct approach can be straightforward to apply.

Introduction

There is considerable debate around discount rates and the treatment of risk in economic analysis. Since the 1980s, most government economic appraisal guidelines in Australia have recommend using a (real) 7 per cent discount rate to account for both the value of a dollar in the future relative to a dollar now and risk (OBPR 2020, p.8). However, support for this position has eroded over time.

Critics have argued that the discount rate should be updated to reflect changing economic conditions, specifically, the reduction in yields on government bonds since the 7 per cent discount rate was introduced (Grattan Institute 2018, p.8). Others have gone further, arguing that the theoretical basis for the 7 per cent discount rate is flawed (Costanza et al. 2021), and that risk should be addressed directly through certainty equivalents rather than indirectly through adjusting the discount rate (BTRE 2005, p.9). Finally, critics including the National Farmers' Federation (2021, p.29) have noted that the discount rate recommended for use in Australia is high relative to other OECD countries. For example, the United Kingdom uses a discount rate which starts at 3.5 per cent and declines over time (OECD 2018).

If these criticisms hold, we are likely to be underestimating the merits of long-term projects versus short term projects and the merits of projects that reduce risk versus projects that increase risk, with implications for the quality of investment decisions. This is especially relevant in agriculture where projects often have long term benefits. For example, a biosecurity program that eradicates a pest that would otherwise become endemic could deliver benefits for generations of Australians.

This report explores both sides of the debate to provide rigorous and practical guidance on how economists should approach discount rates and the treatment of risk in their analysis of agricultural projects. It also includes a real-world biosecurity application to show how the guidance can be implemented and the difference it makes.

We need to discount future costs and benefits

Suppose an agricultural project has $1 of costs in year one and generates $1.05 of expected benefits in year two. Should the project go ahead if the objective is to maximise net benefits? The estimated net benefits would be $0.05 without discounting. This is positive so our initial recommendation would be that the project should happen.

But what if there is an alternative investment that has the same costs in year one and generates higher expected benefits, say $1.07, in year two? We do not want to invest in the agricultural project if it means forgoing a better alternative. This gives the 'opportunity cost' rationale for discounting.[1]

[1] Other rationales for discounting are covered in Boardman et al. (2001, p.227).

The standard approach is to discount based on the expected return on the alternative investment

There are several ways to define the alternative investment (Box 1). Government economic appraisal guidelines tend to assume that the alternative is diversified investment in private assets across the Australian economy, whereas the example above postulates a specific investment. But irrespective of how the alternative investment is defined, the standard approach is to discount based on its expected return. In the example, the expected return on the alternative investment is 7 per cent.[2] Hence, to implement the standard approach, we would apply a 7 per cent discount rate to evaluate the agricultural project (Table 1).

Discount rates convert future values into present values. In this case, $1.05 in year two is worth $0.98 in year one.[3] In other words, we are indifferent between receiving $1.05 in year two and $0.98 in year one. This is because, given an expected return on the alternative investment of 7 per cent, we can invest $0.98 in year one and expect to receive $1.05 in year two.

The estimated net benefits would be -$0.02 with discounting. The negative value suggests that even though the project generates a positive expected return of 5 per cent, we can get an even better expected return from investing elsewhere. On this evidence, our recommendation would be that the project should not go ahead. But is this necessarily the correct recommendation?

---

[2] This happens to broadly align with the average long run return on private investment in Australia of 8 per cent (Harrison 2010, p.59). However, the expected return used for the example could be different and, although the maths would change, the argument being made would be unaffected.

[3] All present values in this paper use year one as the base. Often the first year is called year zero instead. Whether the first year is called year one or year zero does not affect our analysis.

Box 1 What is the alternative investment?

Opportunity cost can relate to either the best alternative investment or the actual alternative investment that would be made without the project. These will be equivalent if the investor is maximising net benefits without constraints. However, they can diverge. For example, governments may be prevented from investing in certain private assets due to competitive neutrality. Even if these assets are the best alternatives, they are not the actual alternatives.

More generally, the alternative investment can be a specific investment or an aggregation over multiple investments. For example, the alternative investment for a horticultural producer could be a specific horticultural project, general investment in the horticultural sector, or investment across the Australian economy.

Table 1 Calculation of net present value for hypothetical agricultural project – standard approach

	Unit	Project values	Discount factors	Project present values
Year one	$	-1	1	-1
Year two	$	1.05	1/1.07	0.98
Sum over years	$	**0.05**		**-0.02**

The standard approach is not always a good approximation

Suppose there are two possible future states of the world – drought and rain. Also assume the agricultural project has zero risk and generates $1.05 of benefits under both possible futures (we relax this assumption later). The alternative investment is higher risk and generates $0.79 of benefits in drought and $1.35 of benefits in rain (Table 2). If both possible futures are equally likely, the expected values are unchanged from the example above.[4]

Table 2 Risk associated with hypothetical agricultural project and alternative investment

	Unit	Project (zero risk)			Alternative (high risk)		
		drought	rain	expected values	drought	rain	expected values
Year one	$	-1	-1	-1	-1	-1	-1
Year two	$	1.05	1.05	1.05	**0.79**	**1.35**	**1.07**

Once we account for risk, it is no longer obvious that the alternative investment is better than the agricultural project. If people are risk averse, they might prefer a certain benefit of $1.05

[4] Expected values are probability weighted averages. For example, the expected benefit for the alternative investment in year two is 0.5 * 0.79 + 0.5 * 1.35 = 1.07.

under the agricultural project to a risky expected benefit of $1.07 under the alternative investment ($1.35 if it rains, $0.79 if it does not).

This can be formalised using certainty equivalents. A certainty equivalent is the minimum amount people would be willing to accept instead of taking a gamble. This provides a risk-adjusted measure of benefits and costs. Box 2 shows how certainty equivalents can be calculated for an investment based on the information provided above and assumptions around income and risk preferences.

Assuming a moderate level of risk aversion, the certainty equivalent for the alternative investment in year two is just $1.03 (Table 3). Comparing the certainty equivalents shows that the agricultural project actually performs better than the alternative investment once we adjust for risk. The recommendation based on the standard approach is not correct in this instance (there are also instances where the standard approach gives a good approximation). This is a potential problem whenever the riskiness of the agricultural project is either substantially higher or lower than the alternative investment it is being compared with.

Table 3 Certainty equivalents associated with hypothetical agricultural project and alternative investment

	Unit	**Project (zero risk)**		**Alternative (high risk)**	
		certainty equivalents	**expected values**	**certainty equivalents**	**expected values**
Year one	$	-1	-1	-1	-1
Year two	$	1.05	1.05	**1.03**	1.07

Box 2 Estimating certainty equivalents for an investment

Certainty equivalents are the best way to account for risk in economic analysis (Stiglitz 2015, p.314). The following approach is based on expected utility theory (BTRE 2005, p.30), which has been used in economics since the 1700s.

1. Simulate incomes over possible futures, with and without the investment.

In practice, these estimates are often based on Monte Carlo simulations. However, alternative stochastic methods are equally applicable. For this box, assume that the 'investment' is the alternative investment above, with benefits of \$0.79 in drought and \$1.35 in rain. This could be the development of a higher yielding variety that responds well to water. (Certainty equivalents can also be calculated for the agricultural project using the same approach.) Suppose that these benefits are received by a farmer whose income without the investment would be \$1.00 in both drought and rain.

With investment:

$$Income_{'drought'} = \$1.00 + \$0.79 = \$1.79$$

$$Income_{'rain'} = \$1.00 + \$1.35 = \$2.35$$

Without investment:

$$Income_{'drought'} = Income_{'rain'} = \$1.00$$

2. Calculate the certainty equivalents, with and without the investment, and take the difference.

2a. Specify a utility function. Commonly used utility functions include the isoelastic utility function (OECD 2018) and the negative exponential utility function (Hone et al. 2020, p.167). Both are single parameter functions that can be parameterised empirically based hypothetical (Hone et al. 2020, p.173) or observed (Groom and Maddison 2019) behaviour in relation to risk. A practical advantage of the isoelastic utility function is that it is dimensionless, making it easier to transfer estimates of risk preferences between studies (see OECD 2018 for a discussion of plausible values). A practical advantage of the negative exponential utility function is that it is defined for zero and negative income, which is often relevant in agriculture. For this box, we are using the negative exponential utility function:

$$Utility = 1 - \exp(-Risk\ aversion * Income)$$

with the risk aversion parameter set to one for simplicity.

2b. Use the utility function to convert incomes to utilities.

With investment:

$$Utility_{'drought'} = 1 - \exp(-\$1.79) = 0.83$$

$$Utility_{'rain'} = 1 - \exp(-\$2.35) = 0.90$$

Without investment:

$$Utility_{'drought'} = Utility_{'rain'} = 1 - \exp(-\$1.00) = 0.63$$

2c. Calculate the certainty equivalents, with and without the investment. The certainty equivalent function can be found by solving the utility function for income. The certainty equivalent function for a negative exponential utility function is:

$$Certainty\ equivilent = \frac{-ln(1 - Expected\ utility)}{Risk\ aversion}$$

With investment:

$$Certainty\ equivilent = -ln\left(1 - \frac{(0.83 + 0.90)}{2}\right) = \$2.03$$

Without investment:

$$Certainty\ equivilent = -ln\left(1 - \frac{(0.63 + 0.63)}{2}\right) = \$1.00$$

2d. Take the difference in certainty equivalents, with and without the investment. This gives a certainty equivalent benefit of $1.03. This is the benefit of the investment to the farmer after adjusting for risk.

Systematic risk can be important

In this example, the farmer is not exposed to drought risk in the absence of the investment. Suppose instead that their income without the investment was $0.50 in drought and $1.50 in rain. The investment would exacerbate their existing drought risk since investment benefits are highest in rain when their income would otherwise be highest. As a result of these interactions, the certainty equivalent for the investment would fall substantially from $1.03 to $0.91. This can be confirmed by repeating the steps above with the updated income assumptions. This shows the potential importance of considering interactions between risks. See BTRE (2005, p.47) for a graphical discussion of certainty equivalents and systematic risk.

Adjusting discount rates to account for project risk also has drawbacks

We need to account for the risk associated with the agricultural project and the alternative investment. Some economists argue that this can be achieved by adjusting the discount rate to account for project risk, potentially as part of sensitivity analysis (Appendix A). The aim is to match the agricultural project with alternative investments of similar risk. Hence, lower (higher) risk agricultural projects would be matched with lower (higher) risk alternative investments, which tend to have lower (higher) expected returns. Other economists disagree with this approach. According to Stiglitz (2015, p.314),

> 'The most common mistake in trying to cope with the uncertainties of the benefits and costs of a project is to argue that in the face of risk, the government should use a higher discount rate'.

There are two theoretical drawbacks with adjusting discount rates to account for project risk. First, suppose we are trying to estimate the costs of a dam wall collapsing in the future. This could present significant risk to people downstream of the dam. If they are risk averse, the certainty equivalent cost will exceed the expected cost, so we need to increase the present value costs.[5] Yet if we attempt to account for the risk of a dam wall collapsing by increasing the discount rate, our adjustment will decrease the present value costs. The adjustment will be in the wrong direction. This is a potential issue whenever people are individually exposed to substantial negative consequences as a result of the project (BTRE 2005, p.7).

[5] There are times where it is appropriate to adjust present value costs downwards because of systematic risk (see Harrison 2010, p.46).

Second, increasing discount rates to account for project risk implicitly assumes that risk grows linearly over time. Figure 1 shows the implicit adjustments for risk over time, after adjusting for the time value of money. This is based on a 7 per cent discount rate, although similar issues arise for other discount rates. The adjustment for risk is initially small, just 7 per cent in year 2. The adjustment for risk reaches 50 per by year 12, and 80 per cent by year 36. The important point is that these adjustments are fixed – they are not linked to actual project risk in year 2, or year 12, or year 36. This can result in substantial bias when the evolution of project risk is nonlinear (BTRE 2005, p.9).

Figure 1 Implicit adjustments for risk over time under the standard approach

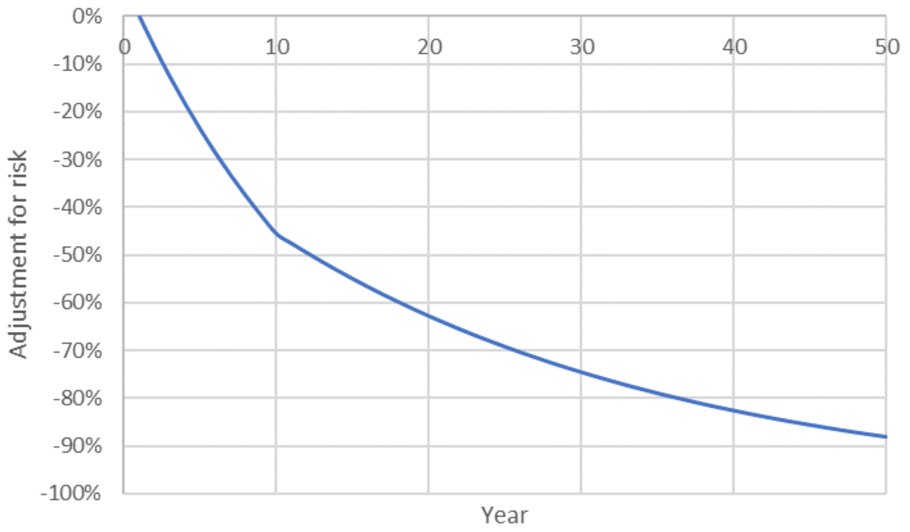

Note: Assumes a constant 7 per cent discount rate under the standard approach and a risk-free discount rate of 0 per cent for the first 10 years and 3 per cent thereafter (Box 3). See Appendix B for a worked example showing how the implicit adjustments for risk were derived.

It is prudent to consider alternatives

We have established that there are technical problems with the standard approach and the most obvious alternative, adjusting the discount rate to account for project risk. These technical problems might be irrelevant for a given project, in which case these approaches would provide a good approximation. However, we cannot know this in advance. And given the sensitivity of the results to discount rates and the treatment of risk, inappropriately applying these approaches could have serious adverse consequences for decision making.

A better approach is to address the time value of money and risk separately

The technical problems described above can be resolved by addressing the time value of money and risk separately. The theoretically correct approach is to use the (real) risk-free discount rate to account for the time value of money and, where risk is important, use certainty equivalents to account for risk.

Use the risk-free discount rate for zero risk projects

The agricultural project described above has zero risk – the costs and benefits are known with certainty. For zero risk projects, discount the project values based on the risk-free return on the alternative investment (Box 3).[6] If the risk-free return on the alternative investment is 3 per cent, this increases the estimated net benefits to $0.02 (Table 4), up from -$0.02 using the standard approach (Table 1). On this evidence, our recommendation would now be that the project should proceed.

Table 4 Calculation of net present value for hypothetical agricultural project – theoretically correct approach for zero risk projects

	Unit	Project values	Risk-free discount factors	Project present values
Year one	$	-1	1	-1
Year two	$	1.05	1/1.03	1.02
Sum over years	$	0.05		**0.02**

[6] It is also appropriate to use this approach for projects where it can be demonstrated that the risk to any individual is small and there is no systematic risk (Arrow and Lind 1970).

Box 3 Estimating the risk-free discount rate

The risk-free discount rate should equal the risk-free return on the alternative investment. Where the alternative investment is not specified, low-risk assets such as government bonds can be used instead. Harrison (2010, p.120) estimates that real returns on government bonds averaged 3 to 4 per cent between 1986 and 2007. However, real returns on government bonds have since fallen. In 2020-21, the average nominal yield on 10-year government bonds was 1.2 per cent, whereas the average forecast break-even 10-year inflation rate was 1.8 per cent (RBA 2021, Table F2.1 and Table G3). Hence, the real return is likely to be close to zero. This has precedent with returns on government bonds being negative in the 1970s.

This means that while a 3 per cent risk-free discount rate may be appropriate in the long run, at the time of writing in November 2021 it is likely to substantially overstate the opportunity cost in the short run. To better capture prevailing conditions, a simple alternative would be to set the short-run discount rate equal to:

$$max\left[\frac{1 + nominal\ yield}{1 + forcast\ inflation} - 1, 0\right]$$

where $nominal\ yield$ is the average nominal yield on 10-year government bonds for the previous financial year and $forcast\ inflation$ is the average forecast break-even 10-year inflation rate for the previous financial year.

Based on the values above:

$$\frac{1 + 0.012}{1 + 0.018} - 1 = -0.6\%$$

Hence, the appropriate short-run discount rate in November 2021 would be zero. The short-run discount rate would apply for the first 10 years. This is the period for which we have reliable market estimates. Thereafter, the discount rate would revert to the long-run average of 3 per cent.

The resulting discount rates are still conservatively high

One reason why our discount rates are conservatively high is that we have not accounted for uncertainty surrounding future risk-free returns. Suppose that risk-free returns could be 1 or 5 per cent with equal probability (and, to keep the maths simple, this is constant over time). The expected present value of a dollar in year 50 will be:

$$0.5 * \frac{1}{(1 + 0.01)^{50}} + 0.5 * \frac{1}{(1 + 0.05)^{50}} = \$0.35$$

This differs from the present value of a dollar using the expected risk-free return of 3 per cent:

$$\frac{1}{(1 + 0.03)^{50}} = \$0.23$$

In effect, the discount rate based on the expected risk-free return is too high, and we need to reduce the discount rate to 2.1 per cent to get the right answer:

$$\frac{1}{(1 + r)^{50}} = \$0.35$$

$$dr = \frac{1}{\$0.35}^{1/50} - 1 = 0.021$$

This is the idea behind Weitzman (1998), who showed that when there is uncertainty over future risk-free returns, the appropriate discount rate is generally lower than the expected risk-free return.

Use certainty equivalents and the risk-free discount rate for projects that could increase or decrease risk

Most projects have the potential to increase or decrease risk – the costs and benefits are not known with certainty. For these projects, estimate the certainty equivalents for the agricultural project in each year (Box 2). Then discount the certainty equivalents based on the risk-free return on the alternative investment. This is identical to the approach for zero risk projects, except that certainty equivalents replace known project values.

This approach is demonstrated in Table 5 with a new hypothetical agricultural project that has the same expected value as before ($1.05) but a certainty equivalent of just $0.90, indicating that the project is quite risky. On this evidence, our recommendation would be that the project should not go ahead. In this case, the standard approach would give the correct recommendation, although it would substantially understate the magnitude of net costs associated with the project (as per Table 1 with a present value of $-0.02).

Table 5 Calculation of net present value for hypothetical agricultural project – theoretically correct approach for risky projects

	Unit	Project certainty equivalents	Risk-free discount factors	Project present values
Year one	$	-1	1	-1
Year two	$	0.90	1/1.03	0.87
Sum over years	$	-0.10		**-0.13**

The theoretically correct approach is straightforward to implement and widely applicable

Implementing the theoretically correct approach is generally straightforward, particularly where stochastic models are already widely used. In this case, it is just a matter of ensuring that income is adequately captured in the stochastic model and adding a few lines to the code or a few columns to the spreadsheet (see example in Appendix B).

The theoretically correct approach is widely applicable. It can be used to evaluate government projects as well as other scenarios, such as the costs of a biosecurity outbreak. The theoretically correct approach was developed for 'normative' applications, including as cost benefit analysis, where the objective is to better understand the effects of agricultural projects on people's wellbeing. However, it also provides a defensible starting point for 'positive' applications, where the objective is to better understand people's behaviour. That said, there are a myriad of ways in which people make decisions in relation to discount rates and the treatment of risk, so alternative heuristic approaches should also be considered for positive applications.

Appendix A: Sensitivity analysis of discount rates

Most Australian economic appraisal guidelines recommend a 'central case' discount rate of 7 per cent. Recognising that this includes a risk premium that might not be appropriate for some projects, they typically also recommend sensitivity analysis with discount rates of 3 per cent and 10 per cent (for example, OBPR 2020, p.8). There are several limitations associated with this approach.

First, the true net present value might fall outside the range of estimated net present values. As discussed above, the current risk-free discount rate is close to zero. Hence, even 3 per cent includes a risk premium. Some projects can reduce risk by essentially providing insurance against negative outcomes, such as drought or pests. Applying a 3 per cent discount rate could substantially underestimate the net present values of these projects.

Second, even when the true net present value falls inside the range of estimated net present values, guidance is rarely provided on the appropriate net present value to use in decision making. Instead, decision makers tend to default to the 'central case' discount rate of 7 per cent.

Third, and more fundamentally, including a risk premium in the discount rate conflates the time value of money and risk (see adjusting discount rates to account for project risk also has drawbacks).

Appendix B: Example biosecurity application

To demonstrate the theoretically correct approach, we have revisited a previous biosecurity application. In 2018, ABARES undertook modelling for the Australian Chief Plant Protection Office on the costs of citrus canker becoming endemic in Australia. The original modelling used the standard approach. We have updated it to implement the theoretically correct approach.

The objective of the earlier modelling was to help the Australian government better understand the benefits of investments to prevent the arrival or spread of citrus canker (in terms of avoided costs). This is somewhat unusual in that the analysis was an input to further analysis, rather than a complete analysis. However, the approach demonstrated below is equally applicable to more conventional investment appraisal problems. The main difference will be in the scenarios modelled. The scenarios for this application are without and with citrus canker. The scenarios for an investment appraisal problem are with and without the investment. In some cases, it might also be necessary to disaggregate the analysis. For example, if an agricultural investment is publicly funded it could be worth modelling the impacts on farmers and other Australians separately.

Citrus canker

As background, citrus canker is a highly contagious bacterial disease that reduces the marketable yield of citrus trees (Figure 2). Citrus canker is widespread in many tropical and subtropical citrus growing areas throughout the world and has been detected and successfully eradicated on several occasions in Australia.

Figure 2 Citrus canker

Original citrus canker model

The original citrus canker model has both biophysical and economic modules. The biophysical module predicts the spread of citrus canker considering local expansion and intermittent long-range jumps. It also predicts where citrus canker would affect agricultural production based on moisture and temperature (see threshold area in Figure 3) as well as the effects on yields, with and without control measures to mitigate the impacts. The economic module predicts the effects of citrus canker on prices received due to the closure of sensitive export markets to Australian citrus. To estimate grower revenue, the area of citrus is multiplied by yield and price, with yield and price being adjusted for citrus canker as required. Grower costs are then subtracted to give grower income (Figure 4).

Figure 3 Predicted extent of citrus canker effects on agricultural production

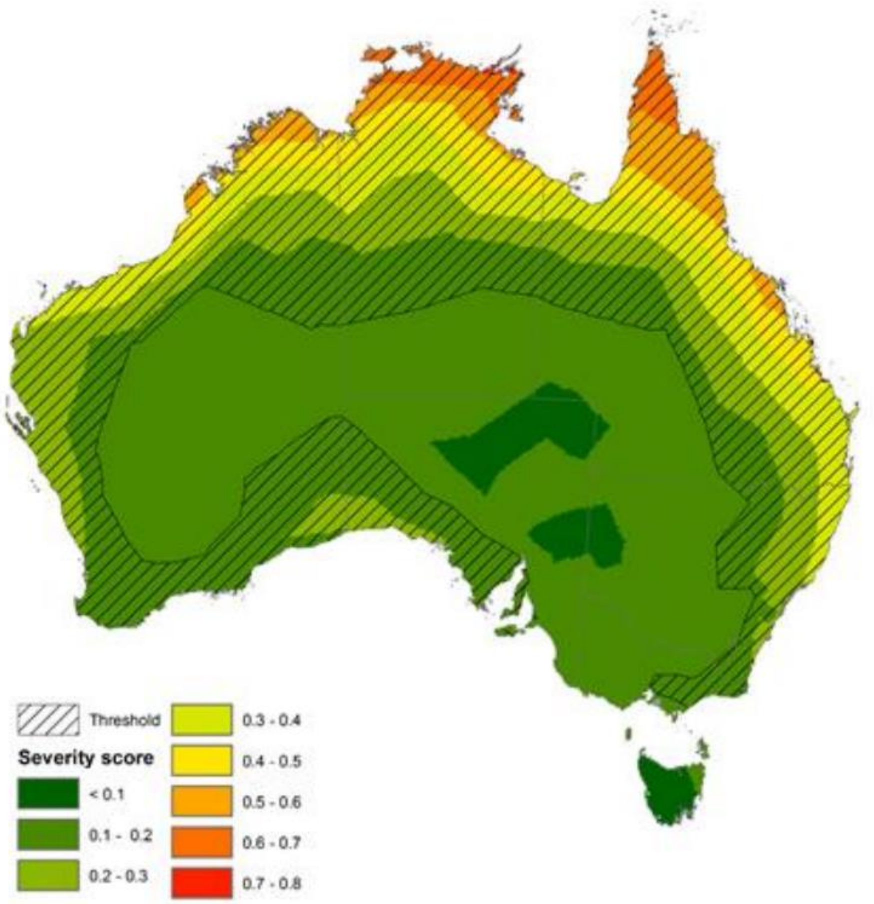

Figure 4 Citrus canker model

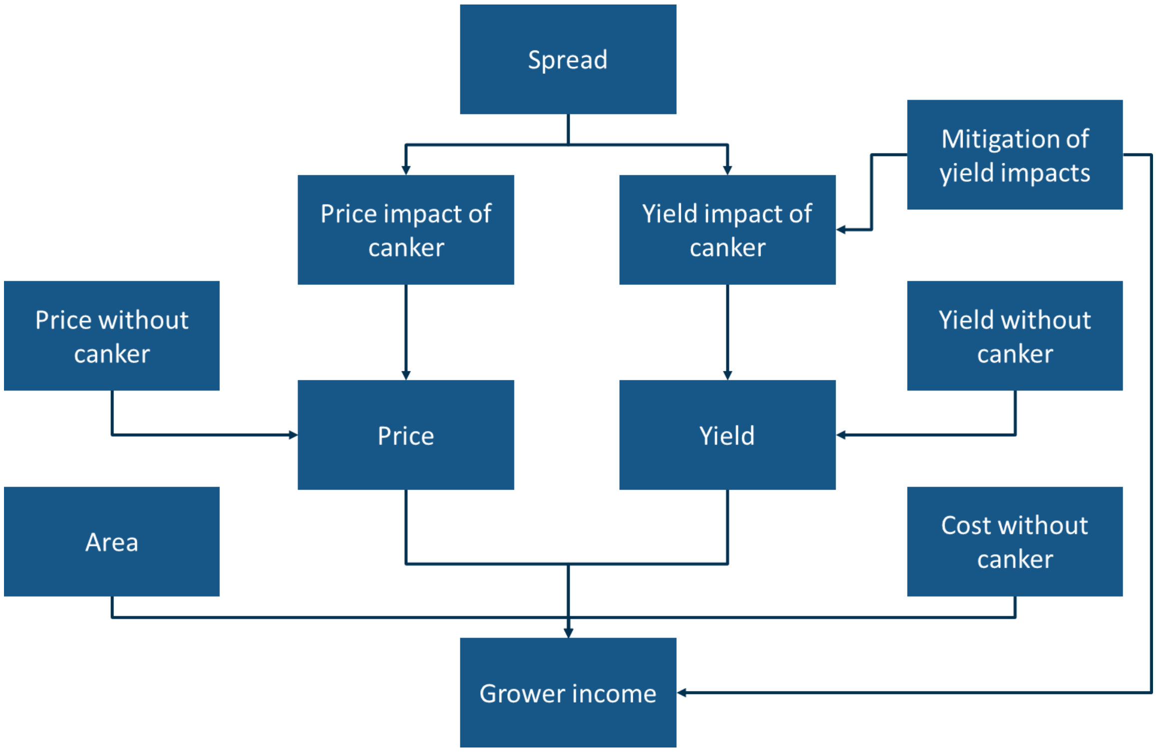

Implementing the theoretically correct approach

We expanded the original citrus canker model to apply the theoretically correct approach. To keep the reporting simple, this appendix only reports the effects on Queensland growers.

1. Simulate grower incomes over possible futures and time, without and with citrus canker

To simulate grower income over possible futures, we converted the original citrus canker model into a Monte Carlo simulation model. This involved identifying parameters that are highly uncertain and important in driving the results (Table 6). We then specified probability distributions for these parameters, drawing on various sources, including citrus canker spread modelling, expert judgement, and statistical analysis.[7]

[7] This Monte Carlo simulation model was developed for illustrative purposes. Some of the data used to estimate the probability distributions were not ideal. For example, the regressions were based on aggregate horticultural industry data rather than disaggregated data from citrus growers. In addition, the model does not account for potentially relevant complexities, such as the statistical dependence between yield and price.

Table 6 Selected stochastic parameters in the expanded citrus canker model

Stochastic parameter	Unit	Distribution type	Distribution values	Distribution source
Spread of citrus canker	Array, spread by region and time	Choice	Equal probabilities	Realisations of stochastic biophysical spread model
Yield impact without mitigation	%	Triangular	Lower = 5, central = 15, upper = 25	Expert judgement
Effectiveness of mitigation	%	Triangular	Lower = 0, central = 50, upper = 100	Expert judgement
Yield growth scalar without canker	Index, current year = 100	Normal	Mean = 0.398, standard deviation = 0.185	Linear regression on historical data
Price growth scalar without canker	Index, current year = 100	Normal	Mean = 0.287, standard deviation = 0.154	Linear regression on historical data
Cost growth scalar without canker	Index, current year = 100	Normal	Mean = 1.325, standard deviation = 0.379	Linear regression on historical data

Note: The model also captures random annual shocks to yield, price and cost.

The Monte Carlo simulations were implemented through a simple script that linked to the original citrus canker model. The script loops over 100 simulations (s) and 50 years (t). For each simulation, the script draws values from the probability distributions, plugs these values into the original citrus canker model, and records the income of growers over time. This generates an array of income values, $Income_{s,t}$.

We ran the Monte Carlo simulation model without and with citrus canker. Figure 5 shows that a wide range of incomes are possible in any given year, with the range of incomes increasing over time. It also shows that citrus canker tends to reduce incomes.

Figure 5 Simulated incomes of Queensland growers, without and with citrus canker, 100 simulations and 50 years

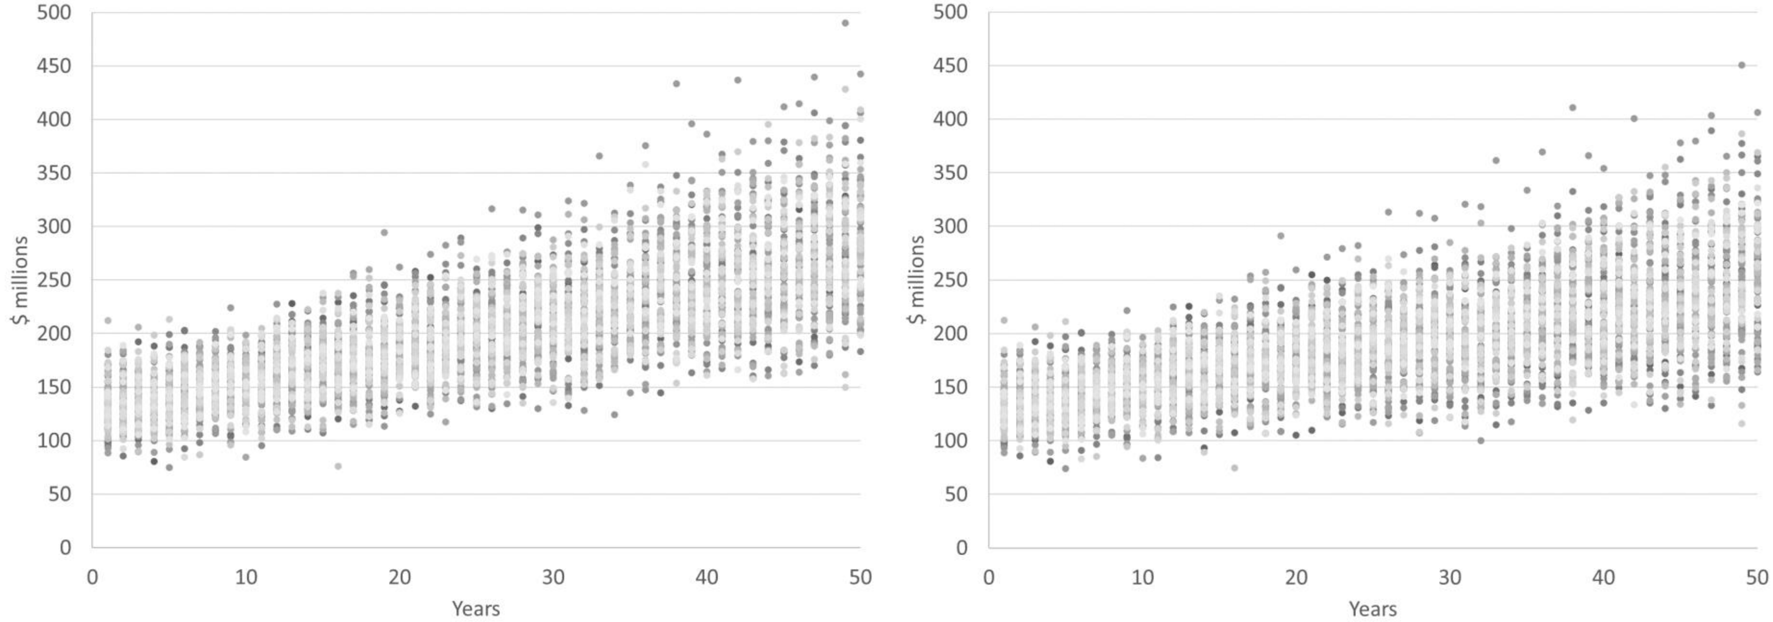

Without citrus canker | With citrus canker

Note: Each grey dot represents a simulation and year.

2. Calculate the certainty equivalents in each time period, without and with citrus canker, and take the difference

The script converts the income arrays into utility arrays using the isoelastic functional form:

$$Utility_{s,t} = \frac{{Income_{s,t}}^{(1-\eta)}}{(1-\eta)}$$

where η is the constant relative risk aversion parameter. The constant relative risk aversion parameter was set to 1.5 based on recent empirical evidence from the UK (Groom and Maddison 2019).[8]

Certainty equivalents are then calculated for each year:

$$Certainty\ equivilent_t = \left[\frac{\sum_s Utility_{s,t}}{Number\ of\ simulations} * (1-\eta)\right]^{1/(1-\eta)}$$

Figure 6 shows the certainty equivalents in each year, without and with citrus canker.

[8] Where the risk aversion parameter has the potential to shift the recommendations of the analysis, drawing an indicative value from the literature may not be sufficient. An alternative is to run a survey to elicit the risk preferences of those affected (for example, Hone et al. 2020, p.173).

Figure 6 Certainty equivilent incomes of Queensland growers, without and with citrus canker

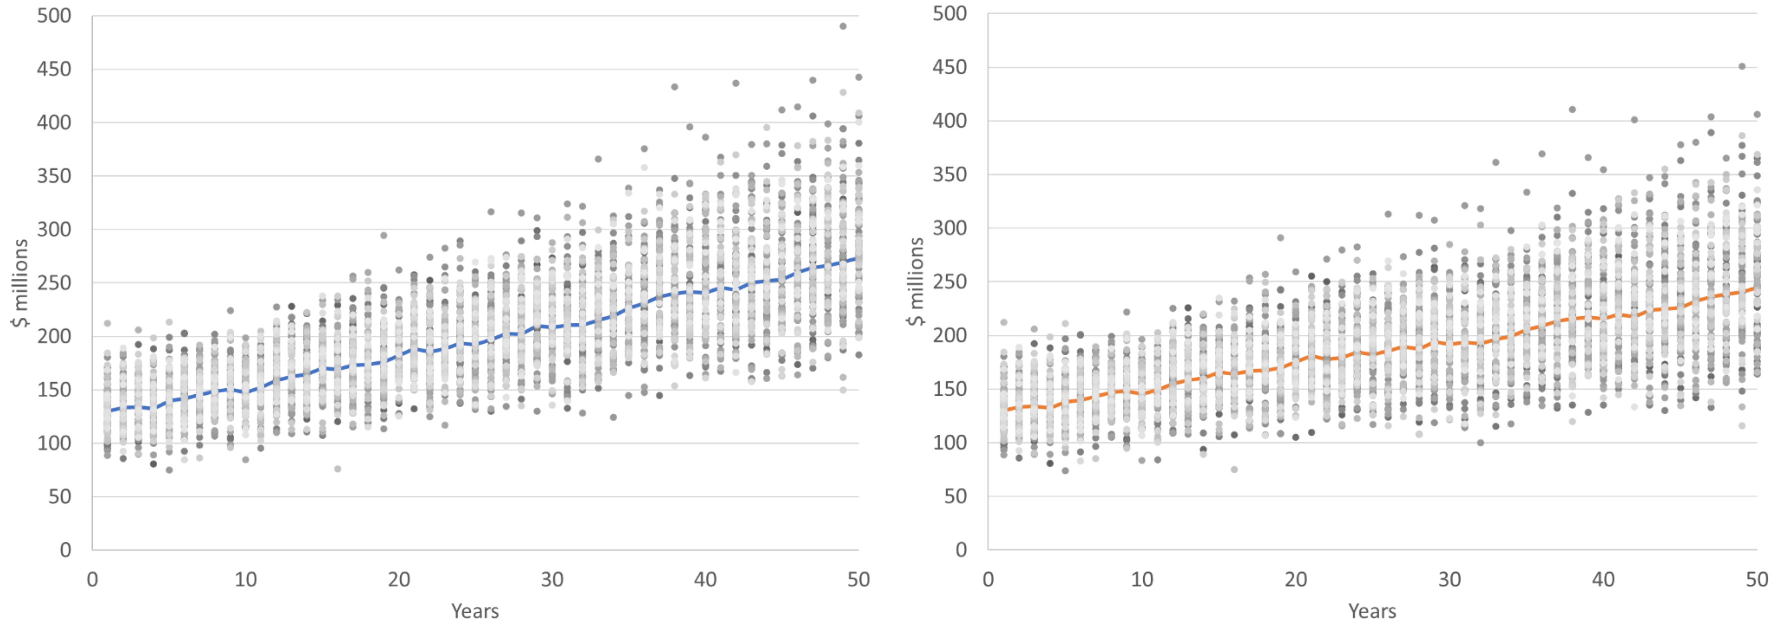

Without citrus canker

With citrus canker

Note: The lines show the certainty equivalents in each year.

And the difference in certainty equivalents is taken:

$$Certainty\ equivalent\ cost\ of\ canker_t = Certainty\ equivalent_t^{without} - Certainty\ equivalent_t^{with}$$

This gives the undiscounted costs of citrus canker to growers, after adjusting for risk. Figure 7 shows that the undiscounted costs of citrus canker increase over time as citrus canker spreads, affecting access to sensitive export markets and citrus yields.

Figure 7 Undiscounted costs of citrus canker to Queensland growers

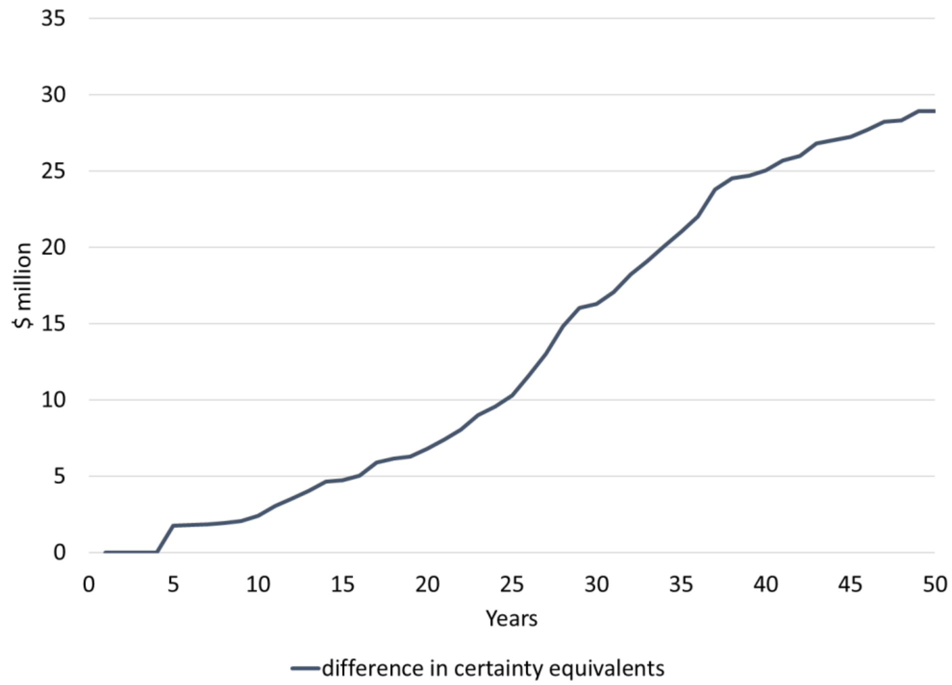

3. Discount using the risk-free discount rate and sum over time

Finally, the script discounts the undiscounted costs using the risk-free discount rate and sums over time:

$$Discount\ factor_t = \begin{cases} \dfrac{1}{(1+srdr)^{t-1}} & for\ t \leq 10 \\ \dfrac{1}{(1+srdr)^9 * (1+lrdr)^{t-10}} & for\ t > 10 \end{cases}$$

$$Present\ value\ cost\ of\ canker = \sum_t Discount\ factor_t * Certainty\ equivilent\ cost\ of\ canker_t$$

where $srdr$ is the short run discount rate and $lrdr$ is the long run discount rate. Consistent with Box 3, we used a short run discount rate of 0 per cent and a long run discount rate of 3 per cent.

Figure 8 shows the effect of discounting. Overall, citrus canker is estimated to cost Queensland growers $320 million in present value terms over the next 50 years.

Figure 8 Discounted costs of citrus canker to Queensland growers

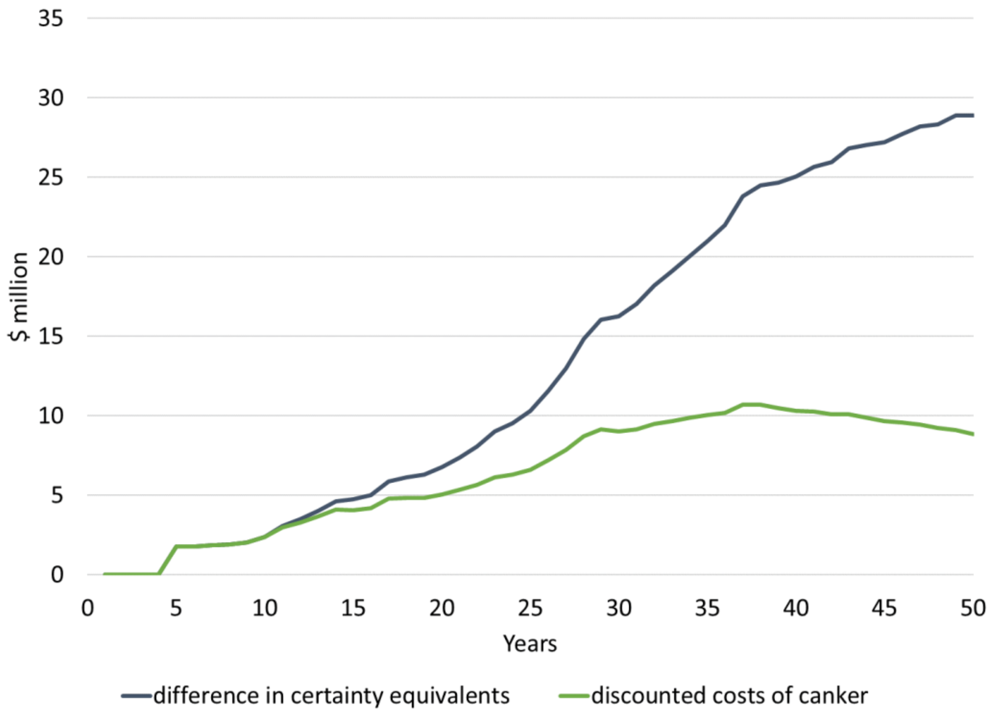

Comparison with the standard approach

By contrast, under the standard approach, citrus canker is estimated cost Queensland growers just $80 million in present value terms over the next 50 years.[9] This large difference is mostly explained by the treatment of risk. For example, in year 25, the theoretically correct approach adjusts for risk by increasing the costs of citrus canker by 6 per cent, whereas the standard approach implicitly adjusts for risk by reducing the costs of citrus canker by 69 per cent. Why do the magnitudes and directions of the adjustments differ?

As discussed above, the theoretically correct approach considers the actual risks associated with citrus canker to growers. Figure 9 presents the estimated income distributions of growers in year 25, without and with citrus canker. As expected, citrus canker reduces mean income across the simulations, from $196 million to $186 million. If we were not adjusting for risk, this difference ($10 million) would be our estimate of the cost of citrus canker. However, citrus canker also increases the standard deviation of income across the simulations, from $31 million to $32 million. That is, citrus canker increases the risks that growers are exposed to. Given that most growers are somewhat risk averse, this increases the costs of citrus canker to growers, albeit slightly. The theoretically correct approach accounts for this by adding $0.6 million to the costs (an increase of 6 per cent). Mechanically, this adjustment happens through the certainty equivalents.

[9] The results presented here differ from the 2018 report because of updates to yield impact and growth assumptions.

Figure 9 Income distrubutions of Queensland growers in year 25, without and with citrus canker

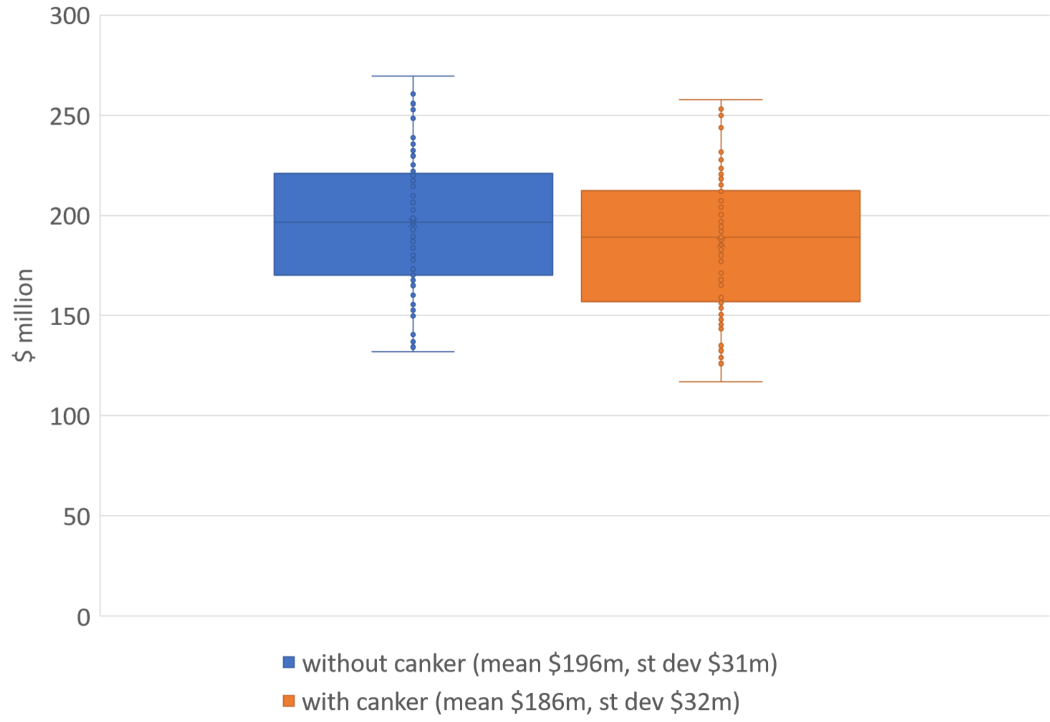

The standard approach is very different. Risk is addressed implicitly through discount factors rather than certainty equivalents. The discount factor in year 25 under the standard approach is:

$$Discount\ factor_{t=25}^{with\ risk} = \frac{1}{(1+0.07)^{25-1}} = 0.20$$

This includes an adjustment for risk. We can estimate what the discount factor would be without the adjustment for risk as follows:

$$Discount\ factor_{t=25}^{without\ risk} = \frac{1}{(1+0.00)^{9} * (1+0.03)^{25-10}} = 0.64$$

Hence, the adjustment for risk is:

$$Adjustment\ for\ risk_{t=25} = \frac{0.20-0.64}{0.64} = -69\%$$

The higher discount rate reduces rather than increases the present value of future costs. So, the direction of adjustment under the standard approach is wrong. But even if the adjustment was in the right direction, there is no reason to think that the magnitude would be right. The problem is the adjustment does not account for the risks associated with citrus canker to growers. It is determined entirely by the year and a generic discount rate. As such, there is no reason to think that the adjustment will be appropriate.

Lessons from biosecurity application

The application demonstrates that the standard approach underestimates the costs by hundreds of millions of dollars in the case of citrus canker. This can have real world consequences, including failing to make worthwhile investments to limit the arrival or spread of citrus canker because we are substantially underestimating the avoided costs from such investments. However, this is just one application and does not mean that the standard approach will always produce misleading results.

The application also shows that the theoretically correct approach can be straightforward to apply. Expanding the citrus canker model to implement the theoretically correct approach took a few hours and only required basic economics and programming. We could have built a more sophisticated stochastic model, but even the relatively simple stochastic model described above was sufficient to show that risk should increase rather than decrease costs, and that the increase should be small.

References

Arrow, L & Lind, R 1970, 'Uncertainty and the evaluation of public investment decisions', *American Economic Review*, vol. 60, no.3, pp. 364-378.

Boardman, A, Greenberg, D, Vining, A & Weimer, D 2001, *Cost-benefit analysis: Concepts and practice*, Prentice Hall, New Jersey.

Bureau of Transport and Regional Economics (BTRE) 2005, 'Risk in cost-benefit analysis', Australian Government, Canberra.

Costanza, R, Kubiszewski, I, Stoeckl, N & Kompas, T 2021, 'Pluralistic discount recognizing different capital contributions: An example estimating the net present value of global ecosystem services', *Ecological Economics*, vol. 183.

Grattan Institute 2018, *Unfreezing discount rates: Transport infrastructure for tomorrow*, Grattan Institute, Melbourne.

Groom, B & Maddison, D 2018, 'New estimates of the elasticity of marginal utility for the UK', *Environmental and Resource Economics*, vol. 72, pp. 1155-1182.

Harrison, M 2010, *Valuing the future: The social discount rate in cost-benefit analysis,* Productivity Commission Visiting Researcher Paper.

Hone, S, Story, J, Von Gerhardt, R, Moore, J, Gormly, R & Fargher, W 2020, *Insurance in the agricultural sector: On-farm financial risk management project*, Aither, Melbourne.

National Farmers Federation (NFF) 2021, *Regionalisation agenda*, National Farmers Federation, Canberra.

OECD 2018, *Cost-benefit analysis and the environment: Further developments and policy use*, OECD, Paris.

Office of Best Practice Regulation (OBPR) 2020, *Cost-benefit analysis: Guidance note*, Australian Government, Canberra.

Stiglitz, J & Rosengard, J 2015, *Economics of the public sector*, W. W. Norton & Company, New York.

Weitzman, M 1998, 'Why the far-distant future should be discounted at its lowest possible rate', *Journal of Environmental Economics and Management*, vol. 36, pp. 201-208.