ASSESSING THE VALUE OF IMPROVED INFORMATION AND MANAGEMENT STRATEGIES FOR OPTIMAL IRRIGATION SCHEDULING

A THESIS SUBMITTED TO THE UNIVERSITY OF MANCHESTER FOR THE DEGREE OF DOCTOR OF PHILOSOPHY IN THE FACULTY OF SCIENCE AND ENGINEERING

2022

THOMAS KELLY

DEPARTMENT OF MECHANICAL, AEROSPACE AND CIVIL ENGINEERING

CONTENTS

LIST OF FIGURES

FIGURES FROM APPENDICES

FIGURES FROM SUPPORTING INFORMATION

LIST OF TABLES

TABLES FROM SUPPORTING INFORMATION

LIST OF DEFINITIONS

- Water productivity: The amount of crop yield produced per unit of water applied through rainfall or irrigation

- Total available water in the root zone (TAW): The quantity of water available for uptake by the crop roots. Calculated by subtracting the soil-water content at Wilting Point from the soil-water content at Field Capacity.

LIST OF PUBLICATIONS

- **Kelly, T.D.** and Foster, T., 2021. AquaCrop-OSPy: Bridging the gap between research and practice in crop-water modeling. *Agricultural Water Management*, *254*, p.106976.

- **Kelly, T.D**., Foster, T., Schultz, D.M. and Mieno, T., 2021. The effect of soil-moisture uncertainty on irrigation water use and farm profits. *Advances in Water Resources*, *154*, p.103982.

- **Kelly, T.D**., Foster, T., Schultz, D.M. 2022. Assessing the value of adapting irrigation strategies within the season. *Agricultural Water Management*, **(in review)**

- **Kelly, T.D**., Foster, T., Schultz, D.M. 2022. Assessing the value of Deep Reinforcement Learning for irrigation scheduling. *Computers and Electronics in Agriculture*, **(in preparation)**

ABSTRACT

Agriculture is the main sectoral user of water globally. Increasing pressures on freshwater resources, coupled with population growth and climate change, mean there is a need to improve agricultural water productivity to tackle water scarcity and food insecurity. Improving irrigation scheduling practices is a key solution to these challenges, in particular in high-productivity agricultural regions where potential gains from improving irrigation application technologies have largely been exhausted.

This thesis consists of four journal articles (two published, one in review and one work in progress) that together form a body of research focused on identifying and evaluating solutions for improving agricultural water productivity through improved irrigation scheduling practices. Work presented in the thesis focuses primarily on case studies in the central United States, where agriculture is the main sectoral user of water, and where there are growing issues of water scarcity related to intensive abstractions for irrigation. Analyses with the thesis leverage novel crop-water modelling tools, including AquaCrop-OSPy that was developed directly as part of this thesis (Chapter 2)

The first empirical paper of the thesis (Chapter 3) examines how uncertainty in the knowledge of soil texture and local weather conditions impacts the water productivity of irrigation decisions made by farmers. This analysis demonstrates that a farmer's choice of irrigation management rules has a much greater impact on productivity and profitability of water use than uncertainty in soil-moisture monitoring. A key conclusion from this paper is that perfect soil-moisture information is not required to make near-optimal irrigation decisions, and that greater benefits can be obtained by improving baseline irrigation management rules and heuristics.

Building off of this finding, the second empirical paper of the thesis (Chapter 4) then examines how scheduling rules can be best adapted to improve the productivity and profitability of irrigation decisions under weather uncertainty. By comparing optimised irrigation heuristics that are applied in every year, to heuristics that are re-optimised during the season, an assessment of the added value of the adaptive in-season irrigation scheduling approaches was conducted. Results demonstrate that robustly optimised fixed management rules can achieve the vast majority of the profits obtainable with perfect

foresight, and that the value of within-season adaptation is generally marginal unless perfect weekly weather forecasts are available to support in-season scheduling adaptation.

Finally, the last empirical paper of the thesis (Chapter 5) explores the potential for AI approaches such as Deep Reinforcement Learning to enhance value of adaptive in-season irrigation scheduling in complex decision-making environments. An assessment is made of the added value of using Deep Reinforcement Learning for irrigation scheduling compared to optimised soil-moisture thresholds. These results show that only under restrictive water caps and minimised weather variability does Deep Reinforcement Learning increase profits compared to optimised soil-moisture thresholds.

DECLARATION

No portion of the work referred to this Thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning

COPYRIGHT STATEMENT

ACKNOWLEDGMENTS

First and foremost I have to thank Tim Foster. I truly could not have asked for a better supervisor and it is because of you that I now consider myself to be a scientist. You gave me the encouragement and confidence to go off in interesting directions, but also knew how bring me back and turn that into publishable research. Your knowledge, patience, and constant support is very much appreciated. None of this would be possible without you.

Next I have to say a massive thank you to David Schultz. You will forever be the gold standard for scientific communication, whether that be writing papers, talks or posters. Your expertise and infectious positivity towards these areas has improved my work immensely.

I am also incredibly grateful to my parents and grandparents for the constant support throughout this process, and for every sacrifice they have made in helping me get to this point.

Thank you to all of the DTP guys for making the past 3.5 years a fun and enjoyable experience. It really wouldn't have been the same without all of the trips to Revs, away days, activities and of course 5-a-side Football.

A big thank you also to my brother, sister, home friends and uni friends for all their support and all the great times we have had.

I also have to thank Cliff and Scragg for being the best friends and housemates I could have asked for. Every pint, game of fifa, and hot fuzz quote has made the past 9 years in Manchester truly great.

Thank you also to Steve, Rachel and Ollie for welcoming me into your family 3 years ago and for all the kindness and support you have shown me.

And finally, the biggest possible thank you to Emily. A week into this PhD I met my soulmate and now my fiancé. You have made not just the past three and a half years incredible, but also all the years to come. Every long walk, activity, day trip, and holiday has created unforgettable memories. I wish you the best of luck finishing your PhD, and hope I can replicate the love and support you have shown me since the day we met.

1 INTRODUCTION

This chapter will start by giving a brief overview of the importance of irrigation for agriculture (Section 1.1), followed then by a summary of different options for improving agricultural water productivity as a response to water scarcity and insecurity (Section 1.2). Key gaps within this research, in particular around solutions for improving irrigation scheduling, will then be highlighted (Section 1.3). Finally, specific research aims for each Chapter will be outlined along with the structure of the thesis (Section 1.4).

1.1 MOTIVATION

Water is a critical input for agricultural production and food security. Globally, insufficient or variable rainfall is considered to be the biggest constraint to crop production (Vadez et al., 2013; Salman et al., 2021). Over 40% of global croplands experience green water scarcity (i.e., local rainfall is not enough to meet crop requirements) for at least 5 months of the year (Rosa et al., 2020). By acting as a buffer against this rainfall variability, irrigation enables farmers to achieve larger and more consistent crop yields, while also enabling land to be cultivated for longer periods of the year in some regions (Howell, 2001). The importance of irrigation for agriculture is illustrated by the fact that irrigated croplands make up only 23% of all global cropland, yet produce 40% of global food supplies (Rosa et al., 2020).

The substantial difference in the productivity of rainfed and irrigated food production makes irrigation an essential tool for combating global food insecurity. However, it also means that agriculture is frequently at the heart of issues of water insecurity and conflict in many parts of the world (Mekonnen & Hoekstra, 2016). Globally, agriculture accounts for over 60% of water withdrawals and over 80% of global consumptive freshwater use (Siebert et al., 2015; Rosa et al., 2020), meaning agriculture is both a driver and a victim of water scarcity issues. Two examples of this driver–victim relationship come from Spain (Expósito & Berbel, 2019; Julio Berbel et al., 2019) and the central United States (Scanlon et al., 2012; Foster et al., 2015; Butler et al., 2020), where excessive groundwater abstraction resulting from irrigation expansion has caused declines in aquifer water storage, which in turn threaten sustainability of agricultural production systems and rural economies. As a consequence of this growing

water scarcity, local policy makers have imposed policy restrictions on irrigation water use, forcing farmers to adopt more water efficient practises, technologies or crop choices.

In parallel with these challenges, it is also widely recognised that food production must increase substantially in the coming decades to meet the demands of growing populations (Howell, 2001; Foley et al., 2011). However, given the ecosystem damage caused by previous agricultural expansion, this increase must take place without further increasing existing pressures on land and water resources (Foley et al., 2011; Zheng et al., 2018). Needing to produce more food with less water has resulted in a shift from optimizing yield and expanding cropping area, towards optimizing productivity on existing agricultural land while minimising use of limited freshwater resources (Howell, 2001; García-Vila et al., 2009; Adeyemi et al., 2017). Achieving these goals is challenging, in particular in the context of increasing crop water requirements (Turral et al., 2010; Kahil et al., 2015) and more volatile rainfall patterns caused by climate change (IPCC, 2022). Increased water productivity via rising carbon dioxide (CO_2) concentrations is unlikely to mitigate these negative impacts due to declines in food nutrient density (Zhu et al., 2018; Dong et al., 2018).

It must also be noted that increasing field-level water productivity (i.e., minimizing water losses from deep percolation and surface runoff), can have unintended consequences when viewed at a basin scale (Van der Kooij et al., 2013; Berbel & Mateos, 2014; Grafton et al., 2018). Losses through runoff and deep percolation are not necessarily lost when taking this wider view of the system as these losses contribute to streamflow and aquifer recharge. Losses through evapotranspiration (ET) on the other hand are generally considered lost to the local hydrological system – though they are beneficial losses to the farmers as ET drives crop growth. Examples from Spain, Chile, China and the United States have found farmers have used their water productivity improvements to expand irrigated area or switch to more water intensive crops (Scott et al., 2014; García et al., 2014; Berbel et al., 2019; Zhang et al., 2020). This behaviour maintains or even increasing their total water abstractions, whilst now decreasing streamflow or return flows to the aquifer. This effect is often referred to as the irrigation efficiency paradox, and highlights that improvements to irrigation water productivity alone may not solve water scarcity issues without also being paired with local government policies such as caps on water abstractions (Scott et al., 2014; Grafton et al., 2018).

1.2 PATHWAYS TO INCREASING AGRICULTURAL WATER PRODUCTIVITY

Efforts to increase food production while minimising use of limited freshwater resources are fundamentally focused on enhancing agricultural water productivity – i.e. the amount of crop yield produced per unit of water applied through rainfall or irrigation. Multiple pathways exist through which agricultural water productivity can be improved. For example, changes to irrigation technologies can enhance the efficiency of water delivery to crops, minimising non-beneficial losses of water, and hence increasing water productivity (Camp, 1998). Alternatively, scheduling practices and rules can be adapted to ensure that water is only delivered at the times and locations where it is needed, avoiding excessive water losses or yield declines through water stress (Adeyemi et al., 2017). Additional pathways to improving water productivity include changes to land management practises (Bossio et al., 2010), as well as the development and adoption of new drought-tolerant crop varieties (Nuccio et al., 2018).

Improving irrigation technologies (e.g., switching from gravity systems to pressurised sprinklers or drip irrigation), has been shown to deliver large improvements to water productivity (Sammis, 1980; Camp, 1998; Al-Jamal et al., 2001; Maisiri et al., 2005; Woltering et al., 2011; Gonçalves et al., 2020). However, the potential gains in water productivity by these improvements reside in agricultural regions with low levels of mechanization, and also entail large upfront costs and training to encourage farmers to adopt new technologies (Koundouri et al., 2006; Mwangi & Kariuki, 2015). In contrast, more developed agricultural regions (e.g., United States) have largely made the switch to more efficient irrigation systems. For example, over 67% of irrigated land in the United States is already irrigated using pressurised sprinkler or drip irrigation systems (USDA-NASS, 2018). As a result, in these regions, there are diminishing returns from further increases in application efficiency through technological change.

In contrast, for mechanised farming systems in high-income countries, the biggest gains in water productivity are likely to be found through improving the way irrigation is scheduled and managed at field levels (Zhang et al., 2021). In the United States, the most common pieces of information reported by farmers to schedule irrigation is personal inspection of the crop (77%) and soil (40%). High-tech solutions such as soil-moisture or plant sensors are only used by 12% and 2% of US irrigating farmers respectively. Indeed, in Nebraska, the

state with the largest number of irrigated acres in the United States, use of soil-moisture sensors is still only approximately 30%. These data show that farmers' current irrigation scheduling practises make little use of new data-driven methods and technologies developed to improve water productivity (Adeyemi et al., 2017; Abioye et al., 2020), indicating there are still large potential increases in water productivity and profitability that can be realised by improving irrigation management practises.

Approaches proposed by researchers to improve irrigation scheduling and management take several forms. Assessing the potential benefits of sensor-based irrigation scheduling, researchers have found that water savings of 15-20% – without reducing crop yields – could be achieved by scheduling irrigation with soil-moisture sensors (Kang et al., 2000; Blonquist et al., 2006; Fereres & Soriano, 2006; Zotarelli et al., 2009). Irrigation can also be scheduled according to the current needs of the crop by calculating reference evapotranspiration (Allen et al., 1998; Najafi & Tabatabaei, 2007; Davis & Dukes, 2010; Adeyemi et al., 2017), which estimates the amount of water consumed/lost through crop transpiration and soil-surface evaporation. Irrigation can also be triggered according to visual or measured crop stress indicators supported by in-field sensing (Jones, 2004; Ihuoma & Madramootoo, 2017; Parkash & Singh, 2020), which will provide farmers with direct feedback on when current soil-water content does not meet crop-water demands.

As well as directly incorporating real-time sensor data into irrigation decisions, crop-simulation models can be used to develop and compare optimal irrigation strategies (Jones et al., 2003; Steduto et al., 2009). Using this approach, previous researchers have optimised irrigation schedules or heuristics (e.g., soil-water content thresholds) within crop-simulation models to maximise water productivity (García-Vila et al., 2009; Schütze et al., 2012; Kloss et al., 2014; Linker & Kisekka, 2017). To improve upon these optimised schedules or heuristics, researchers have also developed frameworks where these irrigation strategies are adapted within the season to account for the unfolding weather or future forecasts (Wang & Cai, 2009; Cai et al., 2011; Hejazi et al., 2014; Jamal et al., 2018, 2019; Linker, 2021). Moving beyond optimizing schedules and simple heuristics, AI-based approaches such as Deep Reinforcement Learning (DRL) (Sutton et al., 1998) have been proposed as solution that can automatically adjust decisions based on new information (e.g., weather forecasts), without needing to be re-optimised. Recent studies have found that DRL can be used to learn

complex irrigation strategies, achieving higher profits compared to rule based strategies (Yang et al., 2020; Chen et al., 2021; Alibabaei et al., 2022)

1.3 RESEARCH GAPS

A key premise of existing research focused on improving irrigation scheduling is that providing farmers with more accurate information about crop, soil, and weather conditions will translate into better water management decisions and hence increased crop water productivity (e.g., Soulis et al., 2015; Adeyemi et al., 2017). However, within this body of research, there has been little focus on quantifying how uncertainty in this information impacts water productivity and farm profits. Further, there has been little focus on assessing the relative contributions of improved information accuracy in comparison with improved scheduling strategies and heuristics. The distinction is important, as it determines whether efforts of researchers and policy makers should be towards simply providing farmers with accurate information, or advising them on improved management practises that perform well even in the presence of information uncertainty. The first of these options could be problematic given the large additional costs of investing in variable-rate irrigation systems, multiple sensors, and field mapping (Hedley et al., 2013; Daccache et al., 2015; Haghverdi et al., 2015). Specifically, is this more accurate information worth the investment? Or can near-optimal irrigation decisions be made with less accurate information (e.g., from a low-cost sensor or soil-water balance model)?

A further key obstacle to developing and operationalising improved irrigation scheduling practices is climate variability. Optimal irrigation scheduling and management rules will vary for each year depending on the quantity and distribution of rainfall during the season. Weather conditions are not known with certainty a priori, and therefore, farmers must make irrigation scheduling decisions under weather uncertainty. A potential solution to this issue is to adapt or re-optimise irrigation strategies within the growing season to account for both the weather experienced so far and future forecasts if available (Wang & Cai, 2009; Cai et al., 2011; Hejazi et al., 2014; Jamal et al., 2018, 2019; Linker, 2021). However, there has been only limited research to date that has evaluated how effective and robust such approaches are in improving agricultural water productivity, given realistic assumptions about farmers' foresight of weather conditions. For example, if strategies were simply

optimised at the start of the season and were based on historic weather data, how would returns compare to more complex re-optimization strategies that may be more costly or time consuming for farmers to operationalise? Addressing this question is critical to improve understanding about the extent to which adaptive irrigation scheduling can improve productivity and profitability of water use, and how robust such approaches are to uncertainty in weather and climate variability.

Finally, to date, much of the work on irrigation scheduling has focused on optimizing a fixed schedule or set of simple heuristics such as soil-moisture target levels aligned with major crop growth stages (Wang & Cai, 2009; Cai et al., 2011; Hejazi et al., 2014; Jamal et al., 2018, 2019; Linker, 2021). However, incorporating the increasing amounts of data available to farmers into simple heuristics will require large numbers of complex rules that will be difficult to design and then optimise. A possible solution comes from AI-based approaches to optimization, such as Deep Reinforcement Learning (DRL), whose success within other complex decision-making environments has motivated its potential for irrigation scheduling (Silver et al., 2018; Vinyals et al., 2019; OpenAI et al., 2019; Badia et al., 2020). Recent work has hinted towards additional value of DRL compared to heuristic approaches (Yang et al., 2020; Chen et al., 2021; Alibabaei et al., 2022). However, these studies have not evaluated rigorously how DRL scheduling approaches perform across different weather conditions and production constraints faced by farmers. For example, under high levels of seasonal weather variability and restrictive limits on water use, how much added value is there in using Deep Reinforcement Learning compared to a set of simple optimised heuristics?

1.4 Research objectives and thesis structure

This thesis presents four articles (two published, one in review and one work in progress) that address key gaps in the research literature on scheduling-based solutions to improving agricultural water productivity described in Section 1.3. The thesis focuses primarily on case studies in the central United States, where agriculture is the main sectoral user of water and where there are growing issues of water scarcity related to intensive abstractions for irrigation (Scanlon et al., 2012; McGuire, 2014; Young et al., 2021).

To support the analysis in Chapters 4 and 5, an open source Python implementation of the crop-simulation model AquaCrop (Steduto et al., 2009) – AquaCrop-OSPy – was developed

to facilitate the development and evaluation of complex and adaptive irrigation strategies. This ability to define any custom irrigation strategy (e.g., customised heuristics or Deep Reinforcement Learning) is not supported in previous AquaCrop implementations or other crop models. This feature, as well as a user friendly interface, straightforward integration with other python libraries, and the capability of running the model in the browser without local installation, have led to over 84k downloads of the model since its release. An accompanying journal article was published to highlight some of the important features and use cases of the AquaCrop-OSPy model, and this article is presented as Chapter 2 of the thesis.

Comparisons between crop, soil, and weather monitoring approaches (e.g., soil-moisture sensors) has often focused on the accuracy of these techniques, with the underlying assumption that more accurate information leads to better decisions. However, an explicit assessment of how the accuracy of such measurements impacts agricultural water productivity or farm profits has not been performed. In response to this gap in the literature, Chapter 3 of this thesis aims to answer the following questions:

1. How does uncertainty in the knowledge of soil texture and local weather conditions impact irrigation water use and farm profits? And hence what is the value of improving the quality of this knowledge?
2. How does that value of improved information compare to the value of improved irrigation strategies?

Even with perfect monitoring of the crop, soil, and current weather, irrigation strategies will always have to be made under uncertainty of future weather conditions. Researchers have often dealt with the issue by re-optimizing irrigation strategies in-season, to account for the unfolding climate or forecasts (Wang & Cai, 2009; Cai et al., 2011; Hejazi et al., 2014; Jamal et al., 2018, 2019; Linker, 2021). Missing from previous works, is a comparison between these adaptive scheduling frameworks with similarly optimised but fixed irrigation heuristics that are optimised to be robust across a range of plausible weather scenarios. To address this gap in the literature, Chapter 4 aims to answer the following questions:

3. What is the economic value of adapting irrigation strategies within the season in response to the unfolding weather?

4. Does this value change depending on adaptation frequency, water-use restrictions, and the presence of perfect short term weather forecasts?

Despite the high performance of optimised heuristics for irrigation scheduling, farmers' real-world decisions must incorporate more information, uncertainty estimation, and general complexity than simple heuristics can represent. Recent successes in Deep Reinforcement Learning in other complex decision-making scenarios has motivated its use for optimizing irrigation strategies (Silver et al., 2018; Vinyals et al., 2019; OpenAI et al., 2019; Badia et al., 2020). Recent publications have suggested this approach can lead to increases in profits compared to heuristic approaches such as single soil-moisture thresholds or fixed schedules (Yang et al., 2020; Chen et al., 2021; Alibabaei et al., 2022). However, as these previous studies evaluated their trained DRL agent against fewer than 3 test seasons, it is unclear how well DRL performs in locations with high year-to-year weather variability. Responding to this gap in the literature, Chapter 5 aims to answer the following questions:

5. Can DRL algorithms learn near-optimal irrigation strategies in the presence of climate variability over a large number of unseen test years?
6. What is the added value of DRL-based irrigation scheduling compared with conventionally optimised heuristics such as soil-moisture thresholds?

Finally, Chapter 6 concludes the thesis, highlighting the main findings from the articles presented, and discusses their implications for researchers, produces, water managers and industry stakeholders. Areas of future work are also identified, as well as indicating how the tools developed in this thesis can assist future research.

2 AQUACROP-OSPY: BRIDGING THE GAP BETWEEN RESEARCH AND PRACTISE IN CROP-WATER MODELLING

This Chapter presents the article:

Kelly, T. D., & Foster, T. (2021). AquaCrop-OSPy: Bridging the gap between research and practice in crop-water modeling. *Agricultural Water Management*, *254*, 106976. https://doi.org/10.1016/j.agwat.2021.106976

The minor changes from the published version include: (1) Figure and Table numbers to match the structure of the thesis. (2) Changes in spelling from US style to UK. (3) Minor change to the final paragraph of 2.3.1 to improve readability.

My contribution to this article was as follows: (1) leading the design of how users would interact with the AquaCrop-OSPy model. (2) All python programming and library development, including converting the model source code from MATLAB into python. (3) Developing the tutorial notebooks and producing the visualizations presented in the article. (4) Writing the content of the article including literature reviews.

AquaCrop-OSPy: Bridging the Gap between Research and Practice in Crop-Water Modelling

T. D. Kelly[1], T. Foster[1]

[1] Department of Mechanical, Aerospace & Civil Engineering, University of Manchester, Manchester, UK.

Corresponding author: Thomas Kelly (Thomas.kelly-2@manchester.ac.uk)

ABSTRACT

Crop-growth models are powerful tools for supporting optimal planning and management of agricultural water use globally. However, use of crop models for this purpose often requires advanced programming expertise and computational resources, limiting the potential uptake in integrated water management research by practitioners such as water managers, policymakers, and irrigation service providers. In this article, we present AquaCrop-OSPy (ACOSP), an open source, Python implementation of the crop-water productivity model AquaCrop. The model provides a user friendly, flexible and computationally efficient solution to support agricultural water management, which can be readily integrated with other Python modules or code bases and run instantly via a web browser using the cloud computing platform Google Colab without the need for local installation. This article describes how to run basic simulations using AquaCrop-OSPy, along with more advanced analyses such as optimizing irrigation schedules and evaluating climate change impacts. Each use case is paired with a Jupyter Notebook, which offer an interactive learning environment for users and can be readily adapted to address a range of common irrigation planning and management challenges faced by researcher, policymakers and businesses in both developed and developing countries (https://github.com/aquacropos/aquacrop).

2.1 INTRODUCTION

Agriculture is the largest sectoral user of water globally, and thus at the heart of debates about how to efficiently and sustainably manage use of limited freshwater resources (Molden & FAO, 2020). Satisfying growing demands for irrigation water as a result of population growth and climate change, while simultaneously ensuring sufficient water remains available to support other human and environmental needs, will require significant improvements in the water use efficiency of food production over the coming decades (Morison et al., 2008; Evans & Sadler, 2008; Adeyemi et al., 2017).

Crop-water models are powerful tools, which can be used to help design and identify strategies to enhance the productivity of water use in crop production and address increasing competition over scarce freshwater resources. Applications include the prediction of climate change impacts on agricultural yields and crop-water demands (Semenov, 2009; Lehmann et al., 2013; Rosenzweig et al., 2014), evaluating the potential effects of irrigation management strategies such as deficit irrigation on production risks (Geerts et al., 2010; Attia et al., 2016), as well as the optimization of irrigation strategies in order to maximise yields or profit while minimizing demands on water supplies from surface and groundwater sources (Wisser et al., 2008; Linker et al., 2018; Foster et al., 2019; Li et al., 2019; Jamal et al., 2019).

While applications of crop-water models in the agricultural water management and planning research are now widespread, the ability of practitioners in government, policy or business to leverage these approaches for real-world irrigation planning and management is limited by a number of factors. For example, implementing more advanced applications of crop-water models, such as optimization of irrigation scheduling rules or heuristics, often requires a high level of programming and data analytics expertise that may not be readily available in end user organizations. This issue is compounded by the tendency of many novel research applications of crop-water models to not provide transparent access to source code necessary to reproduce the specific applications of these models (Kloss et al., 2014; Linker et al., 2018), which themselves are commonly released only as closed-source software packages (Jones et al., 2003; Steduto et al., 2009), further limiting bespoke applications to real-world challenges around panning and management of irrigation water use. Many of the intended users (e.g., irrigation planners, agronomists, farmers, etc.) of crop

models such as AquaCrop cannot be expected to be expert programmers or computer scientists which would be required to reproduce complex coding routines integrating GUI's and optimization algorithms in languages they are unfamiliar with.

In this paper, we present AquaCrop-OSPy (ACOSP), a fully open source Python implementation of the U.N Food and Agriculture Organization's (FAO) widely utilised crop-water productivity model known as AquaCrop (Raes et al., 2009). ACOSP builds on existing implementations of AquaCrop in Matlab/Octave (Foster et al., 2017) and R (Rodriguez & Ober, 2019). ACOSP includes several novel features that address key end user application challenges outlined above, and is up to date with the most recent version of AquaCrop (v6.1) released by FAO (one example comparison is shown in Appendix A, with others available on the ACOSP documentation). Specifically, ACOSP provides an environment to enable a user with even basic Python programming experience to rapidly conduct advanced crop-water management and policy analyses. The growing popularity of Python – fuelled by its flexibility, interoperability, extensive ecosystem and learning resources – has led to Python becoming one of the most popular programming languages globally for environmental simulation, including modelling of agro-hydrological systems. The model is fully self-contained, with a modular design and interactive environment that enables users to flexibly adjust model inputs and integrate with other Python modules (e.g., optimization algorithms). ACOSP source code is accompanied by a series of Jupyter Notebooks (Kluyver et al., 2016) to enable users to learn and develop new applications of the model for different aspects of agricultural water management and planning. The source code – developed using the nbdev framework (https://nbdev.fast.ai) – and notebook tutorials leverage the cloud computing platform Google Colab (https://colab.research.google.com) to enable users to run the model in a web browser without the need to locally install, including via mobile devices that are increasingly being used to support agricultural and irrigation decision-making in both developed and developing countries (Vaishali et al., 2018; Bartlett et al., 2015; Adeyemi et al., 2017). In doing so, ACOSP greatly expands the potential user base and applications of the already widely adopted AquaCrop model, enabling users to flexibly test, adapt and extend the model to address diverse agricultural water management problems globally while still maintaining an interactive learning and application environment through

use of Jupyter notebooks. Users can also freely implement AquaCrop source code and packages outside of the notebook environment in any Python IDE should that be preferred.

In the following sections, we demonstrate a series of key use cases of ACOSP in the context of agricultural water planning and management. First, we summarise how a user can run a model simulation for a selected crop-soil-climate system. Next, we demonstrate a number of important applications of the model, including how a user can apply the model interactively to: (i) evaluate crop-water requirements for a number of irrigation management strategies, (ii) identify optimal (yield maximizing) irrigation strategies given different water use constraints through linkage to Python optimization libraries, and (iii) assessing impacts of climate change scenarios on irrigation water demands and crop yields. Use cases are accompanied by a series of Jupyter Notebooks (available at https://github.com/thomasdkelly/aquacrop), containing more explicit tutorials on how to use the model to reproduce the analysis presented in this article. The reader is encouraged to run these notebooks in combination with this article to better understand the application of ACOSP as a tool for irrigation management and policy. These notebooks can also be used as a basis for development of further customised applications of the code to address local, regional and global agricultural challenges.

2.2 Getting started: Running your first simulation with AquaCrop-OSPy (Notebook 1)

In order to setup and run a simulation using ACOSP, the user must first define a range of parameters and inputs to characterise the cropping system being simulated. These include daily weather time-series, along with parameters specifying the crop, soil, management practices and initial conditions for the simulation. For the vast majority of crop simulation models, these inputs and parameters are defined either interactively through graphical user interfaces (GUIs) or through externally modified input files (Ragab, 2002; Jones et al., 2003; Raes et al., 2009; Foster et al., 2017; Rodriguez & Ober, 2019). However, this approach makes it difficult for models to be integrated into larger code bases or optimization packages. This approach also limits flexibility to conduct large batch simulations of field-scale crop models where many different combinations of input files must be considered

(e.g., to reflect spatial heterogeneity in cropping practices across an agricultural landscape or changes in climate variability over time).

The approach for setting up and running simulations in ACOSP has been designed to overcome these issues, and make it easier to simply and flexibly execute simulations for complex and heterogeneous agricultural production systems. After importing ACOSP into a python editor or notebook, the user is able to interactively define each of the main input types of the model (climate, soil, crop, management practices, and initial soil water and groundwater conditions). In each case, defaults have been predefined such that the number of required user inputs are minimised, but at the same time allowing any of them to be edited by the user as needed. All of these component defaults are specified in the class definitions and can be inspected in the source code. For notebook environments, functions exist that enable the user to directly inspect the source code from the notebook, further streamlining the experience for the user and enabling them to quickly customise simulations without the need for a GUI interface that constrains wider model uses.

Modifying each of these default components is done when creating the component object and full details are shown in Notebook 1. As an example the user could create the default Maize crop by running *maize = CropClass('Maize', PlantingDate='05/01')*, or create a Maize crop with an adjusted Canopy Growth Coefficient by running *maize2 = CropClass('Maize', PlantingDate='05/01', CGC=0.001)*. Similarly for the other model components, users are able to modify the default input classes and develop new customised classes. Figure 2.1 shows an example of the initial setup of crop and soil classes along with initial soil moisture conditions in Notebook 1. The user is able to specify climate input data – the other key requirement by AquaCrop – by either selecting one of the predefined inbuilt climate time series, or by importing external data files using custom helper functions such as *prepare_weather* (Figure 2.1).

AquaCrop-OSPy: Notebook 1.ipynb

File Edit View Insert Runtime Tools Help

+ Code + Text

```
[3] # locate built in weather file
    filepath=get_filepath('tunis_climate.txt')
    # create weather dataframe
    weather_data = prepare_weather(filepath)
```

```
[4] # create sansy loam soil profile
    sandy_loam = SoilClass(soilType='SandyLoam')

    # Create AquaCrop default wheat crop
    wheat = CropClass('Wheat', PlantingDate='10/01')

    # Inital water content set to Field Capacity
    InitWC = InitWCClass(value=['FC'])
```

```
[5] # combine into aquacrop model and specify start and end simulation date
    model = AquaCropModel(SimStartTime=f'{1979}/10/01',
                          SimEndTime=f'{1985}/05/30',
                          wdf=weather_data,
                          Soil=sandy_loam,
                          Crop=wheat,
                          InitWC=InitWC)
```

```
[6] model.initialize() # initilize model
    model.step(till_termination=True) # run model till termination
```

```
[7] model.Outputs.Final.head() # show final outputs
```

	Season	Crop Type	Harvest Date (YYYY/MM/DD)	Harvest Date (Step)	Yield (tonne/ha)	Seasonal irrigation (mm)
0	0	Wheat	1980-03-24	174	8.463380	0.0
1	1	Wheat	1981-03-30	545	8.210405	0.0
2	2	Wheat	1982-03-15	895	8.082336	0.0
3	3	Wheat	1983-03-24	1269	8.382703	0.0

Figure 2.1. The creation of a *SoilClass*, *CropClass* and *InitWCClass* object using built-in soil and crop types, and the initial water content in layer 1 set to Field Capacity. These components are then passed into an *AquaCropModel* which is initialised and run till the end of the simulation, after which, model outputs can then be inspected.

Once the user has defined all input classes and data, these components along with user specified start date and end date for the simulation, are passed into the *AquaCropModel* object (Figure 2.1). The model is then initialised using *model.initialize()*, following which the user can choose to execute the model to run forwards N days using *model.step(N),* or until the end of the specified simulation period using *model.step(till_termination=True)*. These two options provide flexibility to users, allowing planting, harvesting or irrigation decisions to be made dynamically based on either internal model variables or external data. As an example, irrigation decisions can be made on each day from a separate code base which could include groundwater information, weather forecasts, or machine learning models (Sun et al., 2017; Linker & Sylaios, 2016; Adeyemi et al., 2018). This feature enables users to define their own custom irrigation strategies (incorporating any external code or data or management rules), which is either extremely difficult or impossible in other crop-water models that incorporate only a fixed set of pre-defined irrigation management heuristics.

Once the simulation has finished, four different output files will be produced comparable with outputs provided by the original AquaCrop model. The *Flux* output shows daily water flux variables simulated by or inputted to the model, such as total water storage in the crop root zone and precipitation. The *WaterContent* output file provides more granular information about the water storage in each compartment on each day of the simulation period. The *Growth* output reports daily simulated crop growth state variables such as canopy cover and rooting depth. The *Final* output lists the final yield and total irrigation applied for each simulated season. Full details on the output files can be found in Notebook 1, which can also be further customised by users if additional outputs from the model are required.

2.3 IRRIGATION PLANNING AND MANAGEMENT USE CASES

2.3.1 Estimating irrigation water demands under different irrigation strategies (Notebook 2)

A common use case for crop-water models such as ACOSP is estimating irrigation water demands in order to support on-farm and regional scale planning and allocation of limited freshwater resources (Jiang et al., 2016; Araya et al., 2017). One of the main strengths of ACOSP is the ability to rapidly evaluate and compare both simple and complex irrigation strategies, which the user can run as batch simulations specified interactively via the *IrrMngtClass* during model setup.

ACOSP can be used to simulate a range of irrigation strategies such as soil-moisture thresholds, set time intervals, pre-defined calendars and net irrigation that are common to many other crop-water models. ACOSP also contains two features that enables arbitrarily complex irrigation strategies to be defined. First, the model can be run on a daily timestep, and second, any model variables can be extracted or edited before the next day is run. This means that irrigation application decisions can be determined outside the model, using information from weather forecasts, water availability (e.g., hydrological model), or other external databases or models. The chosen depth can then be passed back into the model and applied on the next day. This ability to represent the complex rules, heuristics and behavioural norms that are used by farmers in a given region is not typically available in other crop-water models, and is essential to the accurate estimation of plot-level and regional irrigation water demands to support sustainable and efficient agricultural water management. An example of this feature is shown in Appendix A of Notebook 2 which also showcases examples of more traditional irrigation strategies such as set time intervals and pre-defined irrigation calendars.

To demonstrate the use of ACOSP in this context, we apply the model to evaluate crop yields and irrigation water use over multiple irrigation scheduling strategies. This analysis is done for a case study setup of Maize production in western Nebraska in the United States for the period 1982-2018 summarised in Foster et al., (2015). The strategies contrasted include seasonal soil-moisture thresholds which determine the proportion of total soil water holding capacity at which irrigation is triggered, ranging from 0% to 100%. Constant thresholds such as these are commonly used in modelling studies of irrigation management,

as part of irrigation advisories provided by soil-moisture sensor manufacturers, and by farmers themselves (Blonquist et al., 2006; Gutierrez et al., 2014). This test therefore provides a realistic example of the capabilities of ACOSP to support evaluation of crop yield responses to variable irrigation management decisions.

The results show that in the absence of irrigation, yields are highly variable and often fail completely (Figure 2.2). This dramatic impact of added irrigation illustrates why Nebraska has the largest number of irrigated acres in the United States (USDA-NASS, 2018), reflecting the significant drought risks experienced by many farmers within the state (Wilhelmi & Wilhite, 2002; Hoerling et al., 2014). However, the relationship between soil moisture target strategy, irrigation and crop yields is non-linear, demonstrating a typical pattern of diminishing returns to water beyond approximately a soil moisture target of 40%. These results – and those in Appendix A of Notebook 2 – also show the significant differences in yield and total irrigation between strategies. These differences mean that correctly specifying how farmers make irrigation decisions is essential for accurately estimating irrigation water demands – a challenge given the lack of baseline data on water use in many farming regions worldwide (Montginoul et al., 2016; Closas & Molle, 2018; Foster et al., 2019, 2020).

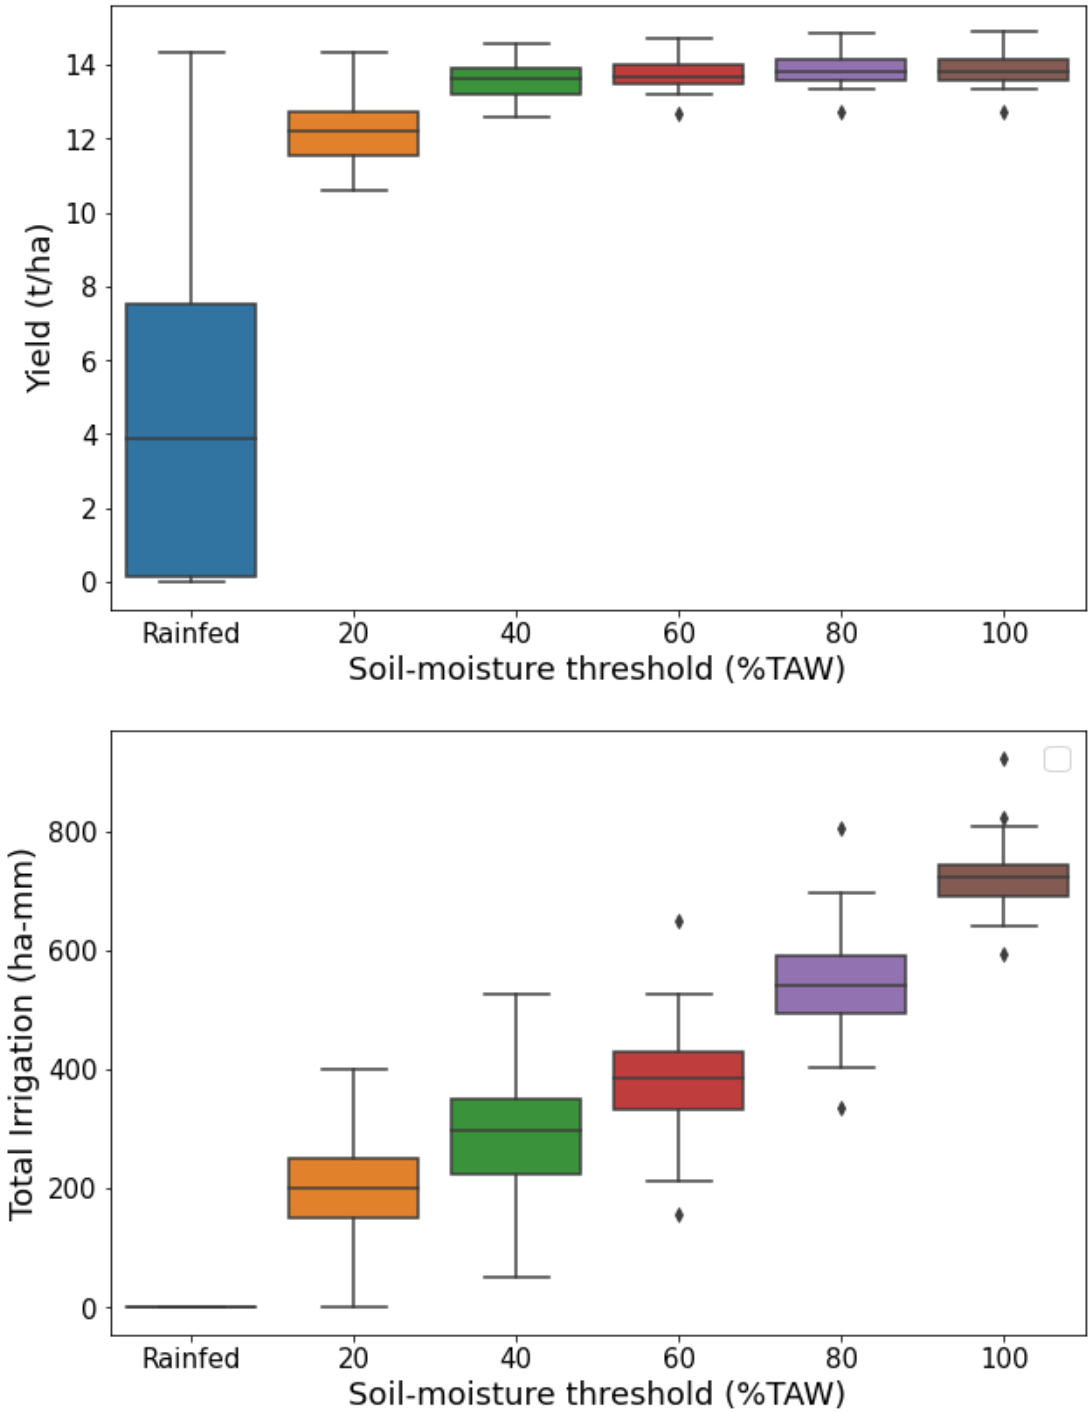

Figure 2.2. A comparison of seasonal (a) yield [tonne/ha], (b) irrigation [ha-mm], for a range of constant soil-moisture thresholds (expressed as a percentage of total soil water holding capacity) over 37 seasons for a maize crop grown in SW Nebraska, United States.

2.3.2 Developing and optimizing irrigation strategies (Notebook 3)

In the previous section, we demonstrated how ACOSP could be used to compare yields and water use for different pre-specified irrigation thresholds. However, it is often the case that farmers and water managers lack information apriori about which irrigation strategy will produce the maximum yield for a given level of water use (Adeyemi et al., 2017). Determining these optimal strategies in practice would require considerable time and investment and so crop-water models are often used in their place. Often this determination is done via brute force search over possible strategies or leveraging optimization algorithms which can be either time consuming or difficult when using a GUI or input file based crop-water models (García-Vila & Fereres, 2012; Araya et al., 2017; Foster & Brozović, 2018; Linker et al., 2018; Jamal et al., 2019).

In this example use case, we show how optimal irrigation schedules can be identified by linking ACOSP with one of the many optimization modules available in the Python ecosystem. Specifically, using the implementation of the model described in Section 2.3.1, we show how ACOSP can be linked to the optimization library *scipy.optimize* (Virtanen et al., 2020) to identify optimal yield-maximizing soil moisture target thresholds over three seasons (2016-2018) for different seasonal water supply constraints (i.e. the maximum amount of water that can be applied in a given season).

The results highlight again the non-linear relationships between crop yields and irrigation water use, along with changes in the distribution of water use within seasons characterised by variations in the optimised soil moisture thresholds (Figure 2.3). In response to increasing water availability the optimal thresholds generally increase (i.e. it becomes optimal to irrigate at lower levels of soil moisture depletion), with water prioritised to mid- and late-season growth stages when overall supply is scarce, as these are the most sensitive periods of crop development for Maize (the example crop used in this Notebook). It should be noted that these optimal thresholds have been found with a local optimiser and so re-running the analysis in Notebook 3 may produce slightly different results than those presented in Figure 2.3.

This example analysis shows how ACOSP can be used to support optimization of irrigation scheduling and water allocation in the context of rising constraints over freshwater availability facing farmers in many regions worldwide (Mekonnen & Hoekstra, 2016). For

example, outputs from this type of analysis can be used to provide advice to farmers about how to effectively optimise water use in the context of hydrologic, technical or regulatory supply constraints, including selection of efficient and profitable intraseasonal irrigation management rules and heuristics. Data can also be used to automate and simplify the development of crop-water production functions such as shown in Figure 2.3, which can be used by policymakers to evaluate expected crop yield responses to different levels of irrigation water supply and make decisions about efficient allocation of water within the agricultural sector and between agriculture and other water users (Igbadun et al., 2007; Foster & Brozović, 2018; Martínez-Romero et al., 2019).

Users could expand this analysis in a number of ways. Our example of ACOSP's optimization integration capabilities was performed for the simple case of a single cropped field with a uniform soil type and uniform weather, however this does not reflect the reality for a farmer irrigating larger and more diverse farms and plots. The example code provided in Notebook 3 could be easily adapted by users to consider more complex optimization problems such as this, treating each individual plot (or sub-plot if considering sub-field level decision making) as an individual instance of ACOSP and optimizing over all plot instances simultaneously. Moreover, as the season unfolds there may be a desire to re-optimise the irrigation thresholds for a given plot or location to better reflect the current climate condition. This again can be implemented in ACOSP simply by saving the current state of the model at day D, and optimizing yield starting from that point using historical data or forecasts. This ability to customise irrigation scheduling rules to accurately reflect real-world decisions, as well as expand the space of potential optimal irrigation strategies, is a key benefit of ACOSP compared to previous crop-water models.

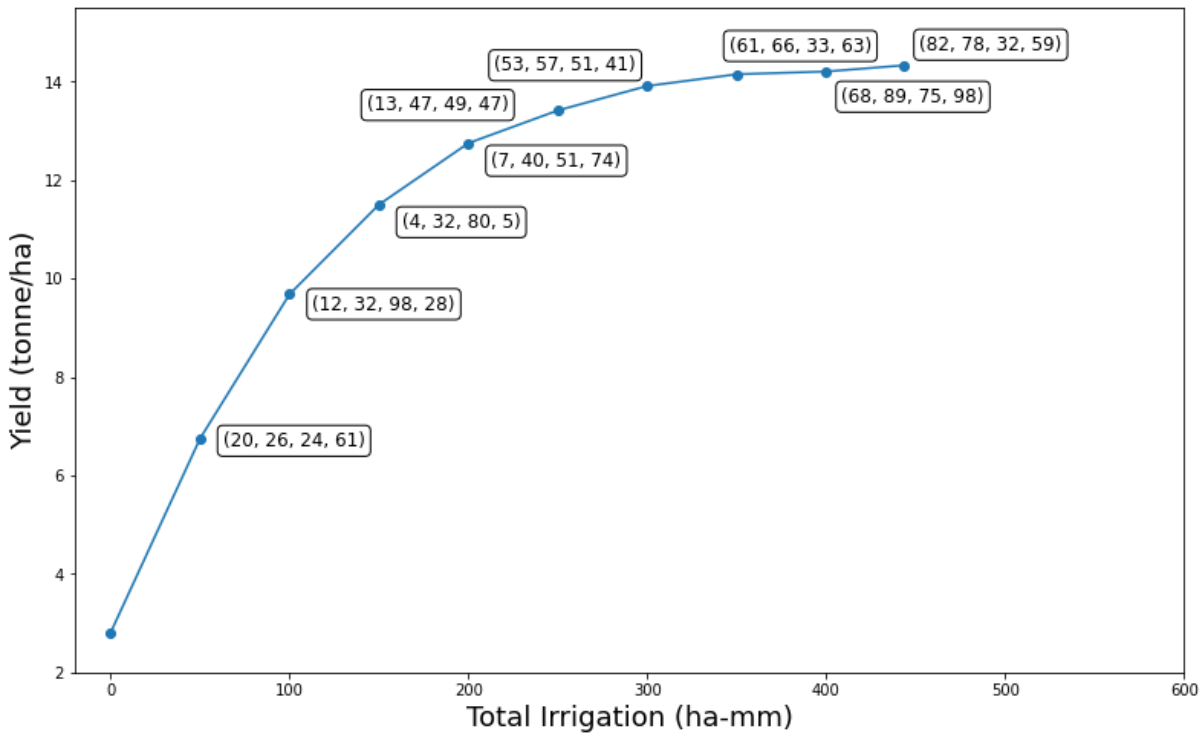

Figure 2.3. Crop water production function showing total maize yield as a function of irrigation water use for an example application in SW Nebraska, United States. The optimization process found irrigation thresholds that maximise yield given a maximum seasonal irrigation limits of [0, 50,…,400, 450 mm]. Optimal irrigation thresholds (expressed as a percentage of total soil water holding capacity) for each water use constraint are also displayed, where each point has four soil moisture target values corresponding each of the main crop growth stages in ACOSP.

2.3.3 Irrigation demands under different climate change scenarios (Notebook 4)

Irrigation is used by farmers to buffer crops against drought and rainfall variability. Climate change is therefore likely to affect irrigation water demands in many regions, and the ability to predict these changes is essential to support sustainable management of agricultural water use and crop production risks (Molden & FAO, 2020). Climate change impact assessments are a common use case for crop-water models as they allow researchers and policy makers to estimate the potential vulnerabilities or future problems for a producing region (Semenov, 2009; Lehmann et al., 2013; Rosenzweig et al., 2014; Dale et al., 2017). These can focus more directly on how climate changes will impact rainfed production, as well as secondary effects such as water use changes for irrigated production. These secondary factors will also impact wider issues such as groundwater depletion (McGuire, 2017) and so it is essential that policy makers and researchers can easily conduct climate change impact assessments using crop-water models like ACOSP.

Building on this need, Notebook 4 guides users through the process of applying ACOSP for a simple analysis of irrigation water demands and crop yields under a range of future climate change scenarios. ACOSP includes inbuilt functionality to directly read and format weather time series for future climate scenarios generated externally using the LARS-WG stochastic weather generator (https://sites.google.com/view/lars-wg) – a widely applied weather generator for agricultural climate impact assessments (Qian et al., 2005; Semenov, 2009; Sha et al., 2019). Conversion of weather data is achieved using the *prepare_lars_weather* function, which includes functionality to calculate reference evapotranspiration inputs required by AquaCrop. Important to note is that as long as the final weather DataFrame passed into the *AquaCropModel* object is in the correct format, the source of the generated data, or the equations used to calculate reference evapotranspiration are arbitrary. This flexibility allows users to generate weather time series offline using alternative weather generators to LARS-WG or other climate datasets/models.

To demonstrate this model application, we use ACOSP to simulate crop yields and irrigation water demands for the same case study of Maize production in western Nebraska considered in the previous notebooks. Our analysis uses weather time series generated by LARS-WG for different future climate time periods (2030, 2050, 2070) and emissions scenario (RCP4.5, RCP8.5) from the EC-Earth climate model (Hazeleger et al., 2012). This

analysis assumes that farmers maintain to the same planting calendar and the same Maize variety, as well as a constant irrigation threshold of 70% total available water throughout all years. These simplifications are made for illustrative purposes of batch climate runs with the model, but could be easily relaxed to consider additional variations in management practices, crop characteristics, etc. to support comprehensive analyses of climate impacts and adaptations on crop yields and irrigation water demands. The results show that yields increase under these future climate scenarios as a result of higher CO_2 concentrations, which translate into higher water productivity within ACOSP (Figure 2.4). Irrigation water use decreases slightly in the climate change scenarios compared to the baseline years, a result that can be explained by the predicted increase in seasonal (i.e. summer) precipitation in Nebraska and the central United States under most future climate change scenarios (Notebook 4).

Users could extend this analysis in many ways, for example by performing a batch analysis over a wide range of climate scenarios, or running ACOSP over a gridded domain to capture spatial heterogeneities or uncertainties (Dale et al., 2017). By linking these gridded model outputs directly to one of Python's many geospatial visualization libraries (e.g. cartopy, geopandas), regional to global climate change impacts on agricultural production and water use can now be easily performed inside a single Jupyter Notebook. Computation efficiency of such analysis can be further improved by parallelizing over multiple processers, with users able to install and run ACOSP on any cloud computing service (e.g. Google Colab). An example of how ACOSP simulations can be run in parallel is shown in Notebook 3.

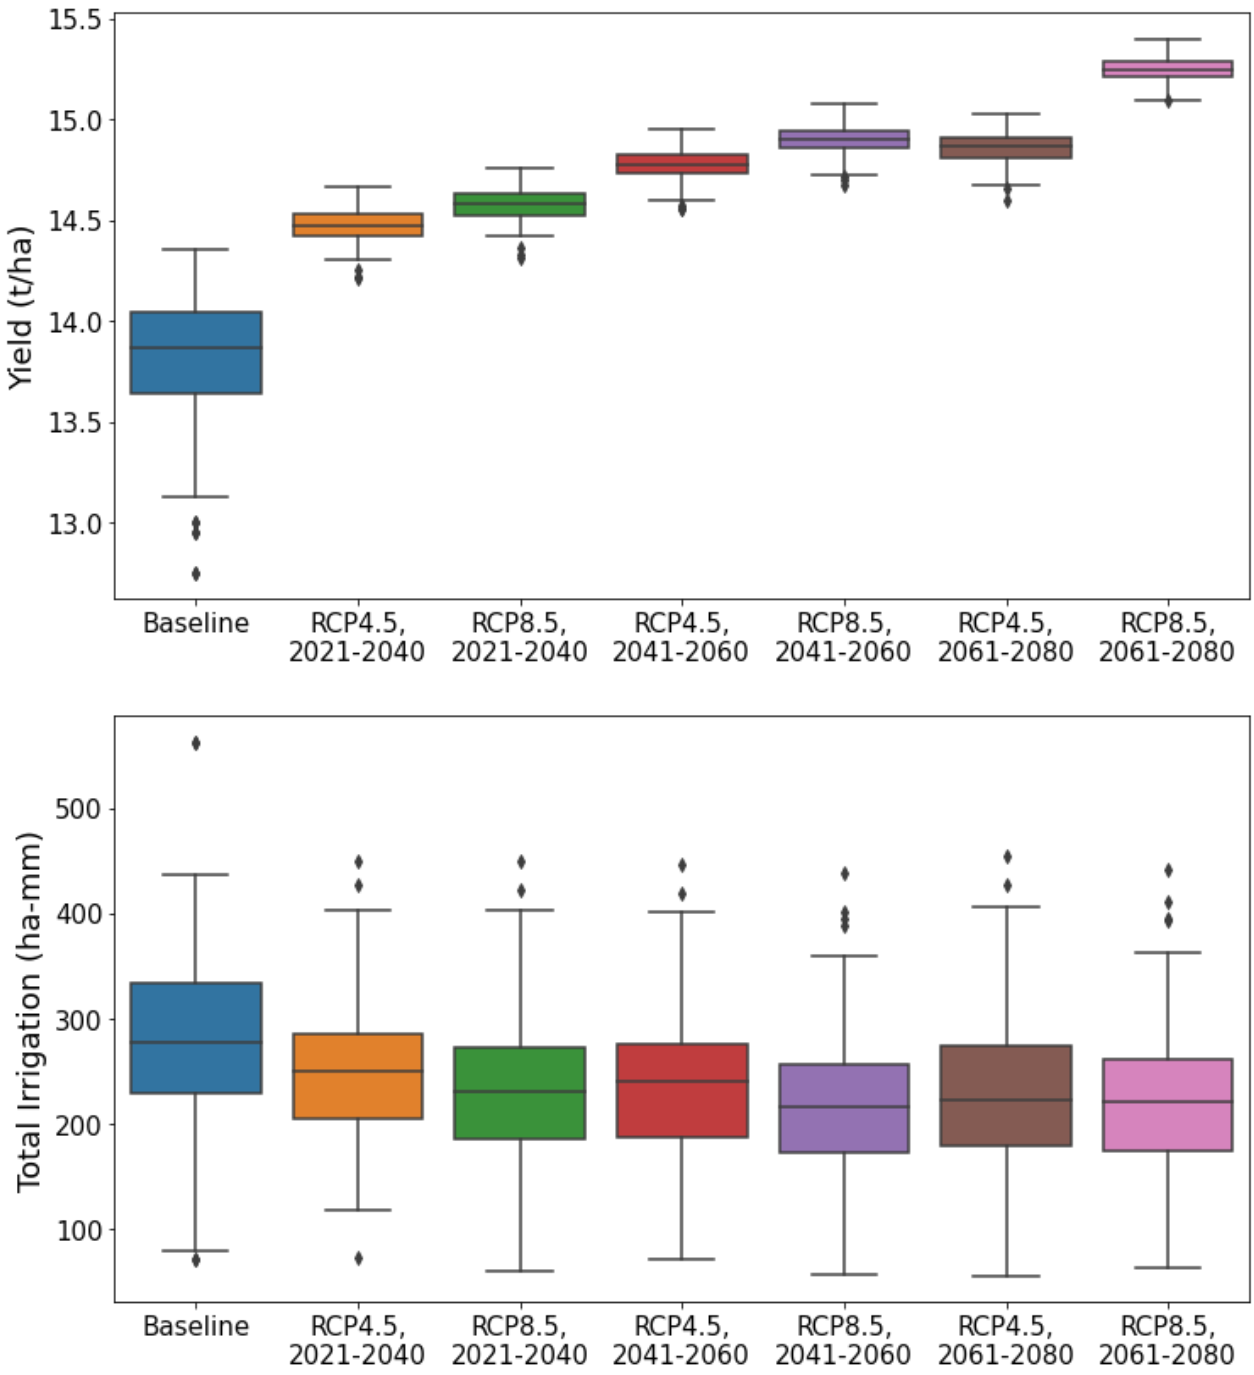

Figure 2.4. Comparison of (a) yields and (b) total irrigation for maize crop in SW Nebraska, United States for baseline climate conditions (1982-2018) and a range of future climate scenarios. Analysis assumes farmers follow a fixed soil moisture target strategy (irrigate once soil water depletion exceeds 30% of soil water holding capacity) in all years and climate scenarios.

2.4 OUTLOOK

AquaCrop-OSPy (ACOSP) is a free and fully open-source tool that allows researchers, industry and policy makers to rapidly leverage crop-water simulation to support irrigation planning and management at farm to regional scales. The model features a user friendly interactive code design, and provides an interface to platforms such as Google Colab to enable a user to quickly conduct simulations via their web browser and leverage cloud computing resources. The extensive Python ecosystem means ACOSP can be easily integrated with the latest machine learning, web development, statistical and data analysis libraries, as well as any other python code base. ACOSP also enables users to bypass the licensing and operating system constrains associated with previous widely used versions of AquaCrop.

Uptake of the model by users conducting applied research and teaching on agricultural water management and policy is supported through development of extensive documentation, including a series of Jupypter Notebooks (https://github.com/thomasdkelly/aquacrop) containing tutorials for a range of common use cases of the model related to irrigation planning, management and climate adaptation. These notebooks are also accompanied by an online forum for users (https://forum.aquacroposforum.com), which will provide a space for the user community to discuss bugs, share example applications and contribute to future development of AquaCrop. For example, opportunity exists for users to contribute to refinement of internal model routines (e.g. more advanced soil moisture flow or nutrient cycle algorithms) or novel applications of the model (e.g. satellite data assimilation and automated parameter calibration) that would be impossible to implement in a scalable or efficient manner with closed-source software.

Through these innovations, we expect that ACOSP will help address a number of the issues that have constrained the use of crop-water models in both policy, practise and applied research surrounding agricultural water management at field, catchment and global scales. Future work will support the maintenance of core model source code in line with future AquaCrop developments, while spurring and enabling new innovations and adaptations to help address global challenges of food and water security facing the agricultural sector.

APPENDIX A: EXAMPLE COMPARISON OF AQUACROP-OSPY WITH FAO GUI AND MATLAB VERSIONS

```
res = run_comparison(model,'tunis_test_1')
```

```
python seasonal mean: 7.52 kg/ha
MAE from windows: 0.02 kg/ha
MAE from matlab:  0.003 kg/ha
```

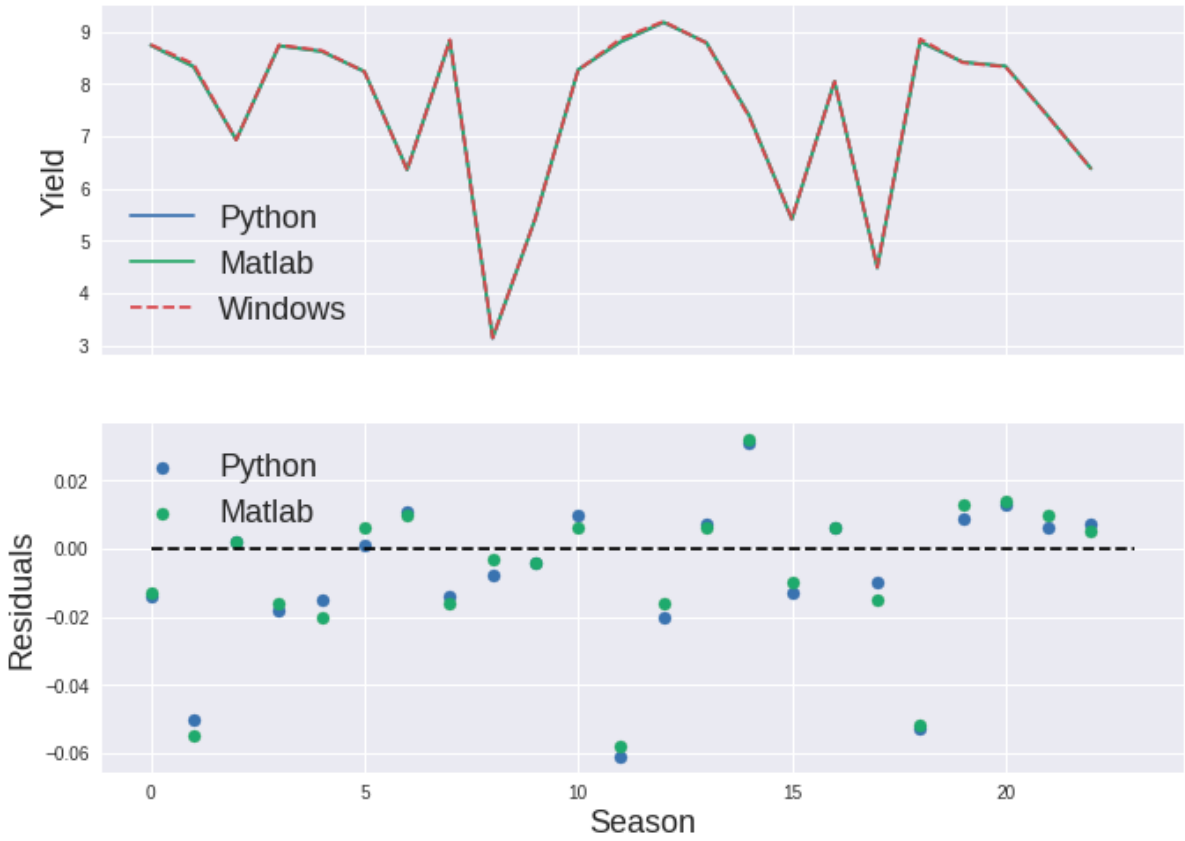

Figure 2.A1. Comparison of yields between AquaCrop-OSPy (python – blue), AquaCrop-OS v6.0 (Matlab – green) and AquaCrop v6.0 (Windows GUI – red dotted) for rainfed Wheat production in Tunis. Also presented is the mean-absolute-error (MAE) in seasonal yields between these versions demonstrating the extremely close agreement between simulated yield outputs from the Python and original FAO (MAE of 0.02 tonne/ha) and Matlab (MAE of 0.003 tonne/ha) model versions. Example taken from AquaCrop reference manual. Other example comparisons can be found on the AquaCrop-OSPy documentation (https://aquacropos.github.io/aquacrop/notebooks/05_comparison)

Acknowledgments, Samples, and Data

Thomas Kelly was supported by the National Environmental Research Council's Understanding the Earth, Atmosphere, and Ocean Doctoral Training Programme, Grant NE/L002469/1.

3 THE EFFECT OF SOIL-MOISTURE UNCERTAINTY ON IRRIGATION WATER USE AND FARM PROFITS

This Chapter presents the article:

Kelly, T. D., Foster, T., Schultz, D. M., & Mieno, T. (2021). The effect of soil-moisture uncertainty on irrigation water use and farm profits. *Advances in Water Resources*, *154*, 103982. https://doi.org/10.1016/j.advwatres.2021.103982

The changes from the published version include: (1) Figure and Table numbers to match the structure of the thesis. (2) Changes in spelling from US style to UK. (3) Additional text at the end of 3.2.4 to clearly summarise the number of experiments as well as computational requirements of the analysis. (4) Minor change to the final paragraph of 3.3.3 and the second paragraph of 3.4.2 to improve readability.

My contribution to the article was as follows: (1) Development of the research questions and methodology. (2) All programming required to run the framework and analyse results. (3) Analysing results and producing tables and figures. (4) Writing the contents of the article including literature reviews.

The Effect of Soil-moisture Uncertainty on Irrigation Water Use and Farm Profits

T. D. Kelly[1], T. Foster[1], David M. Schultz[2, 3], T. Mieno[4]

[1] Department of Mechanical, Aerospace & Civil Engineering, University of Manchester, Manchester, UK.

[2] Department of Earth and Environmental Sciences, University of Manchester, Manchester, UK.

[3] Centre for Crisis Studies and Mitigation, University of Manchester, Manchester, UK.

[4] Department of Agricultural Economics, University of Nebraska-Lincoln, Lincoln, NE 68583, USA

Corresponding author: Thomas Kelly (Thomas.kelly-2@manchester.ac.uk)

ABSTRACT

Technologies that increase the accuracy of soil-moisture monitoring, such as in-situ sensors, have been proposed as a key solution for increasing agricultural water productivity. However, quantifying how uncertainty in soil-moisture estimates lead to irrigation inefficiencies or economic losses has not been explicitly studied. We develop a framework that combines a crop simulation model with a rule-based irrigation decision-making algorithm to assess the impact of soil-moisture uncertainty on irrigation use and farm profits. We apply this modelling framework to a case study of irrigated maize production in Nebraska, United States, a region where improvements in agricultural water productivity are at the forefront of water-policy debates. We consider two main sources of uncertainty that result in a divergence between the farmers' perception of soil-water content and the true water status, namely errors in the knowledge of soil texture and measurement of daily soil-water flux inflows and outflows. Even for very large errors in both soil-texture and water-flux measurements, impacts on water use and profits were marginal (11 ha-mm increase and \$27 ha^{-1} decrease, respectively). In contrast, farmers' choice of irrigation strategy had a much larger impact on water use and profits than uncertainty in soil-moisture information used to implement that strategy. Our findings show that near-optimal irrigation decisions can be made without perfect soil-moisture information. This conclusion suggests that providing farmers with improved irrigation scheduling recommendations – utilizing crop-water models and optimization techniques – would have a larger impact on water-use efficiency than simply providing farmers with technologies to more accurately monitor soil-moisture conditions.

3.1 INTRODUCTION

Increases in irrigation water demand driven by climate change and population growth are expected to exacerbate water scarcity and conflict over limited freshwater resources in many regions around the world (e.g., Molden, 2013). To manage these challenges, there is a need to identify ways to improve the productivity of agricultural water use – commonly referred to as generating "more crop per drop". One key pathway for raising agricultural water productivity is through improvements to irrigation scheduling, which enables farmers to target limited water supply to when the crop is most sensitive to water stress and thus

maximise yield returns to limited freshwater inputs (Morison et al., 2008; Evans & Sadler, 2008; Adeyemi et al., 2017).

The majority of the developments in irrigation scheduling come under the label of *precision irrigation*, which broadly encapsulates solutions that aim to vary the temporal and spatial distribution of irrigation to match crop-water demands (Adeyemi et al., 2017). Associated technologies to support implementation of these strategies include sensors to monitor soil-water content or crop growth (Jones, 2004; Dobriyal et al., 2012; Ihuoma & Madramootoo, 2017), field mapping that highlights heterogeneities in soil texture and other properties (Daccache et al., 2015), and decision-support tools that aid in the timing of irrigation events (McCarthy et al., 2010). The common goal that links these methods is the precise monitoring and management of soil-water content and crop-water needs. Currently, the uptake of precision-irrigation methods and technologies remains low, with most farmers even in developed agricultural systems such as the United States relying on proxies, such as the visual condition of the crop and soil, as well as simple water-balance calculations to determine when to irrigate (Figure 3.1). Much like soil-moisture sensors, these traditional techniques aim to monitor and manage soil-water content in the crop root zone, and hence, the amount of water available for crop uptake. The difference between farmers with and without soil-moisture sensors, therefore, is the accuracy of this monitoring along with the associated rules and heuristics by which this information is used to make irrigation decisions.

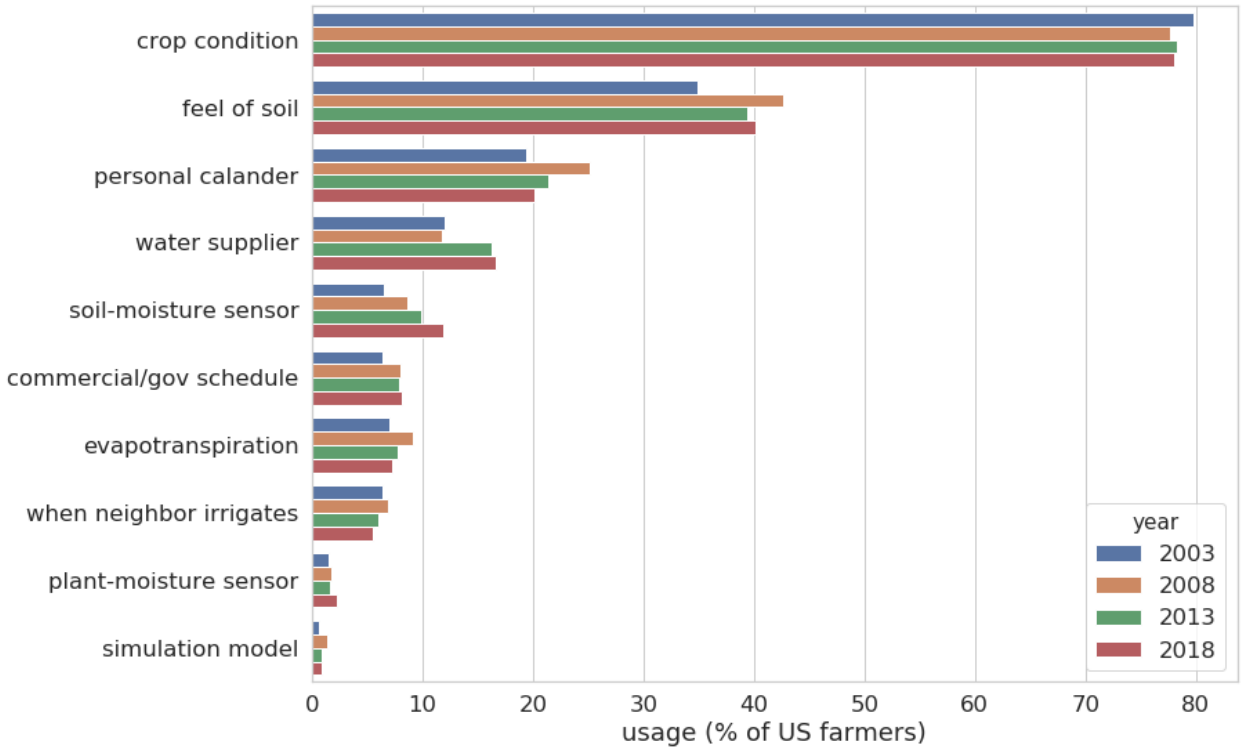

Figure 3.1. Reported methods used by farmers who irrigate in the United States to schedule their irrigation events. Data is from the four most recent 5-yearly (2003–2018) Irrigation and Water Management Surveys conducted across the United States by the Department of Agriculture (USDA-NASS, 2018).

Sources of uncertainty in soil-moisture monitoring and estimation are similar for both traditional and precision-irrigation methods. One important source of uncertainty is heterogeneity in soil textural composition, which affects estimates of current and potential soil-water storage both laterally and vertically within a field (Nielsen et al., 1973; Brocca et al., 2007; Feki et al., 2018). Heterogeneity in soil texture means that point measurements of soil-moisture status are only valid for a specific location and depth, potentially leading to errors when this information is used in water-balance calculations across large areas (e.g., field or farm). Additional uncertainty in soil-moisture estimation also results from errors in climatological (rainfall, evapotranspiration, etc.) and irrigation data characterising inflows and outflows from soil-water storage. Differences in infiltration rates due to differences in soil texture and topography, as well as large distances between the measurement instruments and cropping area, are major contributors to this uncertainty (Chaubey et al., 1999). Moreover, the accuracy of evapotranspiration estimates are greatly affected by the amount of site-specific information used in its calculation (Sentelhas et al., 2010), including in the calculations of reference and actual evapotranspiration (Allen et al., 1998). The result of either soil-texture or water-flux errors is that a farmer may miscalculate the amount of water stored in the soil on any given day, leading to inaccurate estimates of when or how much irrigation should be applied. As a result, a farmer may apply insufficient or excessive amounts of water, with resulting negative impacts on crop yields and/or farm profits.

Precision-irrigation research has focused on developing tools and techniques to reduce the magnitude of soil-moisture error (Soulis et al., 2015; Adeyemi et al., 2017; Kukal et al., 2019). Alongside this, when farmers invest in more accurate soil-water monitoring technology such as sensors or weather stations, they will often also receive advice about the best trigger levels and irrigation management strategies to maximise yields, profits or water efficiency (UNL, 2019). The relative contribution of improved scheduling rules versus improved soil-moisture information to farmers' irrigation water use practices has not been explicitly studied. However, this distinction has important implications for farmers and water managers seeking to understand which interventions and technologies will be most effective for delivering desired improvements in the productivity and value of irrigation water use.

In response to these gaps in the literature, this article presents a framework for studying the impact of varying levels of soil-moisture uncertainty on irrigation water use decisions and farm profits. The framework consists of an optimization module to generate irrigation strategies, a crop-growth model to assess impacts of these irrigation strategies on water use and crop yields, and a mechanism for applying variable soil-moisture measurement errors within model simulations. We apply this framework to a case study of irrigated maize production in Nebraska, United States, using the AquaCrop-OS crop-water model to assess how assumptions about irrigation management practices, and the accuracy of soil moisture information on which these strategies are implemented, affect irrigation decisions and farm profits. Our findings provide insights on the response of irrigation water productivity to soil-moisture measurement uncertainty, showing that the impact of inaccurate soil-water monitoring on crop yields, profits and water use efficiency is low relative to the impacts of using sub-optimal irrigation scheduling heuristics. We discuss the implications of these results for local policy makers, industry and producers, and provide recommendations for how our results can be used to make improvements in agricultural water productivity at field to catchment scales.

3.2 METHODS

In this section, we describe our proposed framework for modelling farmer irrigation decision-making under varying levels of soil-moisture measurement uncertainty, and for assessing the resulting impacts of this measurement uncertainty on water use, crop yields and economic returns. We first present our generalised assessment framework in three main subsections: determination of farmers' irrigation scheduling rules (Section 3.2.1), assessment of the performance of these strategies under different sources and magnitudes of measurement error (Sections 3.2.2), and the incorporation of a feedback mechanism to represent realistic farmer behaviour (Section 3.2.3). We then describe the specific case study application of this framework, including details of model parameters and choice of crop-growth model, optimization algorithm, crop type, and location used as an illustrative example in this study (Section 3.2.4).

3.2.1 Determining farmers' rules in scheduling irrigation

Despite a large variation in the methods used by farmers to schedule irrigation (Figure 3.1), the goal of most methods is to maintain soil-moisture levels above a pre-specified threshold chosen to meet crop-water needs and balance trade-offs between costs and yield benefits of water use. In line with these assumptions and past research (Linker et al., 2016; Foster & Brozović, 2018; Linker et al., 2018), we assume that a farmer's irrigation decision-making can be characterised by the choice of four soil-moisture thresholds throughout the growing season, one for each major crop-growth stage (emergence, early-season canopy development, mid-season crop growth and yield formation, late-season canopy senescence). These four soil-moisture thresholds determine the allowable root-zone soil-moisture content (expressed as a percentage of total soil-water holding capacity) which water contents must fall below before irrigation is triggered. At the point irrigation is triggered, we assume that water will be applied to refill the soil profile back to field capacity, subject to a maximum application rate which is determined by the type of irrigation system and water source used by a farmer.

The choice of soil-moisture thresholds – which we refer to as the *irrigation strategy* – can have a major impact on water use and crop growth. Setting these thresholds too high (i.e., irrigating at low levels of soil-moisture depletion) will result in over-irrigation, whereas setting thresholds too low (i.e., irrigating only at high levels of soil-moisture depletion) will result in under-irrigation and thus production losses. In line with past research showing that farmers' input use decisions may be influenced by their risk preferences (Abdulkadri, 2003; Menapace et al., 2013), we define the optimal set of soil-moisture thresholds as those that maximise a farmers' certainty equivalent (CE). The CE is defined as the amount of economic return that would have to be guaranteed in order for an individual to be indifferent towards a higher but less certain return. The CE has been used in previous agro-economic models of farmer input use decision-making (Lehmann et al., 2013; Foster et al., 2014). In the context of irrigation optimization, the CE for a given irrigation strategy s is defined as

$$CE(s) = E[P(s)] - 0.5 * r * \frac{Var[P(s)]}{E[P(s)]}, \tag{3.1}$$

where strategy $s = [t_1, t_2, t_3, t_4]$ is made up of four soil-moisture thresholds; the expected profits $E[P(s)]$ and variance $Var[P(s)]$ will be calculated over a given set of seasons that characterise farmers' expectations of potential inter-annual weather variability; and *r* is the risk coefficient representing farmer's risk preferences. For risk-averse behaviour (*r* > 0), more weight will be placed on reducing year-to-year variability in profits (Equation 3.1), whereas risk-neutral behaviour (*r* = 0) reduces the right-hand side of Equation 3.1 to expected seasonal profits. For a given irrigation strategy s, the seasonal profit *P(s)* is calculated as

$$P(s) = M * Y(s) - C * I(s) - F, \tag{3.2}$$

where *M* is a constant crop market price [$ per tonne], *Y*(*s*) is crop yield [tonne per ha], *C* is a constant irrigation cost [$ per ha-mm], *I*(*s*) is total seasonal irrigation applied [ha-mm], and F is fixed production costs [$ per ha].

Finding a set of soil-moisture thresholds that maximise CE is a difficult task due to large number of potential soil-moisture threshold combinations and interdependencies between the choice of threshold for each individual growth stage. For example, irrigating more in the early part of the season will affect the optimal irrigation threshold for the rest of the season. As a result, the search space of possible threshold combinations and their resulting outcomes is non-linear and discontinuous. For these types of search spaces, swarm intelligence algorithms have become popular in recent decades (Mavrovouniotis et al., 2017) and have been used previously for the purpose of optimal irrigation planning and management (De Paly & Zell, 2009; Noory et al., 2012). These algorithms operate on the principle that a group of agents scattered across the search space can search locally for solutions while communicating with the other agents. This combination ensures that the solutions found are close to the global optimum.

For this analysis, we determine farmers' optimal (i.e., CE-maximizing) soil-moisture thresholds using the particle-swarm optimization algorithm (Eberhart & Kennedy, 1995;

Freitas et al., 2020), although in principle any optimization method could be used within our framework. This method starts by initializing a population of particles (i.e., possible soil-moisture threshold combinations) across the search space. Each particle is then evaluated and updated using the best combinations found by both the particle and the population as a whole. This evaluate–update process repeats until a stopping criteria has been met (Mathworks, n.d.). Particle-swarm optimization has been successfully applied in a diverse array of research areas, such as control problems, biomedical research, robotics and distribution networks (Poli, 2008).

3.2.2 Quantifying the impact of uncertainty on irrigation decisions and profits

As shown in Figure 3.1 and discussed earlier, many farmers do not measure soil moisture directly using sensors. Instead, most rely on proxy measures of soil-water content, such as visual inspection of the soil and simplified checkbook-style water-balance models. Each of these approaches carry uncertainties and errors in the measurement of soil moisture, which may influence the performance of selected irrigation scheduling strategies. In this study, we focus on two main sources of error: (1) the specification of soil textural characteristics that cause a persistent seasonal bias to water-balance estimation and (2) measurements of daily precipitation, irrigation, and evapotranspiration fluxes that cause daily unbiased errors in soil-moisture content. Our analysis does not consider structural model errors such as simplifications in the water-balance equations inherent in traditional soil-moisture monitoring approaches used by farmers, and our analysis assumes that the underlying soil-water balance model is a 'perfect' model of the chosen system. Focusing on these two sources of error allows us to examine the effects of measurement uncertainty on farmers' irrigation decision-making.

The framework created to study the effect of soil-moisture uncertainty consists of two identical crop-model simulations running in parallel. Simulation 1 represents the farmers' expectations about soil-moisture conditions, which will be impacted by any errors in soil-texture and water-flux measurements. The sources of uncertainty considered are errors in soil-texture inputs added at the beginning of each season and water-flux inputs added on each day of the simulation. Simulation 2 represents the true status of soil-moisture conditions based on the actual soil-textural properties and water fluxes. Irrigation decisions are determined based on the soil-moisture content given by Simulation 1, as well as the

farmers' CE-maximizing irrigation strategy that determines the level of soil moisture at which they intend to trigger irrigation. Irrigation events estimated in Simulation 1 are then applied in Simulation 2 to determine resulting impacts on water use and crop yields. Under perfect information (zero errors in soil texture and water flux), the farmers' perception of soil-moisture conditions match the true state, and both simulations are identical such that the farmer is able to implement perfectly the CE-maximizing irrigation strategy. When the magnitudes of soil-texture errors, water-flux errors, or both are greater than zero, the difference between the farmer's expectations of soil-moisture conditions and reality diverge. This difference potentially results in sub-optimal irrigation scheduling, with negative impacts on crop yields or profits. For example, if the farmer triggers irrigation at a lower level of soil-moisture than the intended CE-maximizing threshold (i.e., because their estimated soil-moisture levels in Simulation 1 over-estimate true soil-moisture levels given in Simulation 2), then the crop may experience unintended water stress and a reduction in yields. Alternatively, if the farmer triggers irrigation at a higher soil-moisture level than the intended CE-maximizing threshold (i.e., because their estimated soil-moisture levels in Simulation 1 under-estimate true soil moisture levels given in Simulation 2), then irrigation water may be wasted and profits reduced.

Two examples of how the added soil-moisture uncertainty causes the two simulations to diverge are shown in Figure 3.2 (where percentage errors drawn from a normal distribution with zero mean and a standard deviation of 30% have been added to the soil-texture and daily water-flux inputs in Simulation 1). The soil-texture errors result in seasonal biases towards either over or under-estimation of drainage and therefore soil-moisture storage, whereas the water-flux error results in random daily errors in the farmer's estimate of the true daily change in water content. When this calculated soil-moisture content (Simulation 1) drops below the pre-defined threshold, irrigation is triggered. This irrigation depth is applied to both fields regardless of whether the true moisture content (Simulation 2) has crossed the intended irrigation threshold. Figure 3.2a shows an example of irrigation being triggered several days after the true water content in Simulation 2 had dropped below the threshold, resulting in water stress and yield losses. In the alternative case, irrigation may be triggered too early, causing profit losses from excess water application and potentially leading to crop damage from temporary waterlogging (Figure 3.2b).

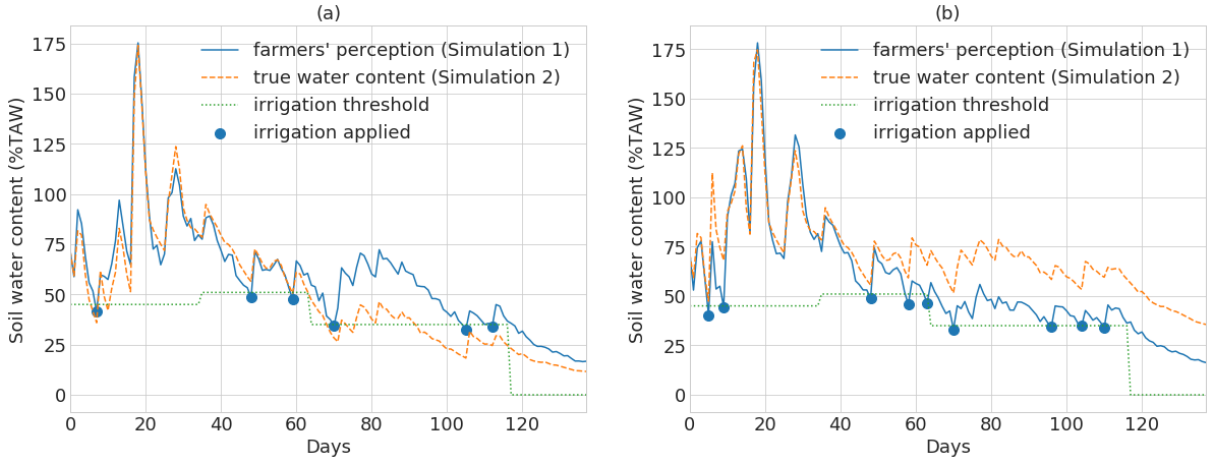

Figure 3.2. Soil-moisture content (% total available water) during an example season where (a) the farmer is overestimating the water content, (b) the farmer is underestimating the water content. A soil-water content of 0% indicates wilting point whereas 100% indicates field capacity. Irrigation events (blue circles) occur when the farmers' perception of soil-moisture content (blue solid line) drops below the irrigation threshold (green dotted line). This irrigation event is applied to both simulations regardless of the true soil-moisture content (orange dashed line). In Figure 3.2b, the added soil-texture and water-flux errors have resulted in an underestimation of the true soil-water content, leading to excess irrigation application and economic losses. This example adds a 30% standard error (meaning percentage errors are drawn from a normal distribution with a mean of zero and a standard deviation of 30%) to the water-flux and soil-texture measurements, which can cause unrealistically high measurement errors such as the irrigation depth error on day 70 of Figure 3.2a.

3.2.3 System Feedback

When incorporating error in farmers' estimation of soil-moisture content, the potential exists for unrealistic irrigation behaviour to occur during simulations that may overstate the true effects of measurement error on farmer irrigation decision-making, crop yields, and profits in reality. For example, the calculated moisture content (Simulation 1) may recommend no irrigation despite the true water content (Simulation 2) being sufficiently low that the crop would be showing visible signs of water stress (e.g., wilting or early senescence). Alternatively, the calculated water content may recommend irrigation despite the true soil profile being visibly saturated or water logged. In these instances, a farmer would likely adjust their irrigation application in response to these clear visual feedbacks from the true field, which would be readily observable through day to day crop management activities even if the farmer was no longer personally measuring the soil-moisture content in the field. This assumption is further supported by the fact that visible inspection of the crop and soil are by far the most common information sources when deciding to schedule irrigation (Figure 3.1).

To capture these behavioural feedbacks, an additional rule is added to our parallel simulation approach to override the daily irrigation recommendation from the farmers' calculation (Simulation 1). This override occurs if, on the previous day, the soil-moisture conditions in the true field (Simulation 2) were either (i) low enough to trigger early canopy senescence, which is an indication of severe crop water stress, or (ii) higher than the soil's true field capacity, which is indicative of visibly moist soil conditions. In these instances, an irrigation event within our assessment framework would be triggered or blocked, respectively. The altered irrigation depth is applied to both fields (Simulations 1 and 2), and the parallel simulations subsequently continue as described in Section 3.2.2. The added feedbacks have the effect of stopping excessive irrigation in the early season, while also preventing unrealistically large water stress in the late season (Figure 3.S1). We hypothesise that incorporating these behavioural feedbacks will minimise the effects of large measurement errors, leading to less importance being placed on the quality of improved soil moisture information – a hypothesis we test in subsequent sections as part of our illustrative case study application.

3.2.4 Illustrative Application

To illustrate our modelling framework, we apply the methods described in Section 3.2.2 and 3.2.3 to the example of centre-pivot irrigated maize production in Nebraska, United States. Maize is the dominant irrigated crop type in Nebraska, which has the largest number of irrigated acres by state in the United States (USDA-NASS, 2018). Irrigation water is sourced primarily from the underlying High Plains Aquifer, and improving productivity of water use is a key priority for policymakers to address growing issues of aquifer depletion (McGuire, 2017), streamflow depletion (Szilagyi, 2000), and degradation of freshwater ecosystems (Palazzo & Brozović, 2014; Perkin et al., 2019). Financial incentives to farmers to encourage adoption of soil-moisture sensors are seen as a key mechanism to improve water productivity and irrigation efficiencies (CropWatch, 2019b), making maize production in Nebraska a logical choice for examining the impact of soil-moisture uncertainty on water use and farm profits.

As the basis for soil-water balance and crop-growth simulations in our illustrative analysis, we use the crop-water model AquaCrop-OS (Foster et al., 2017). AquaCrop-OS is an open-source version of the original AquaCrop (Raes et al., 2009) developed by the Food and Agriculture Organization of the United Nations to simulate crop response to water stress. The model has been successfully applied and calibrated for maize production in a wide variety of environments (Hsiao et al., 2009; Abedinpour et al., 2012), including in our study area in the central United States (Foster et al., 2015) and within the United States and globally (Heng et al., 2009; Foster et al., 2015; Sandhu & Irmak, 2019). AquaCrop-OS captures the impacts of water stress on multiple aspects of crop development (expansion and senescence), transpiration, and yield formation through a series of water-stress thresholds and coefficients that have been calibrated for maize in our study areas as part of previous studies (Forster et al., 2015). These determine the level of soil moisture depletion at which each process becomes constrained by water stress, along with how resulting deficiencies in crop development or transpiration impact on end-of-season biomass production and crop yields. Further details about the simulation of water stress on crop growth and yields in AquaCrop are given in Steduto et al. (2009) and Raes et al. (2009).

The first step in our analysis is to identify CE-maximizing soil-moisture thresholds using particle-swarm optimization as described in Section 3.2.1. Given the non-linear and discontinuous nature of the search space of possible thresholds, the same optimization procedure did not always produce the same set of thresholds. Multiple optimization repetitions were therefore required to capture the full distribution of CE-maximizing strategies. The number of repetitions was set to 50 after observing that the variance between optimal threshold combinations stabilised for more repetitions than 50 (Figure 3.S2). Each of these 50 sets of thresholds can now be thought of as a separate farmer with their own, slightly different, CE-maximizing irrigation strategy. Table 3.1 shows the top five strategies in terms of the resulting 30-year CE.

Subsequently, to evaluate the impact of information uncertainty, each CE-maximizing irrigation strategy was then applied as described in Sections 3.2.2 and 3.2.3. For each of the 50 CE-maximizing irrigation strategies, we conducted a 30-year simulation using our framework's two parallel-model approach. The first model, representing the farmer's expectation of soil-moisture status, perturbs soil-texture properties (% sand and clay content) at the start of the simulation by a normally distributed percentage error (Figure 3.S3). The soil hydraulic properties are then estimated from this soil texture inside AquaCrop via pedotransfer functions (Saxton & Rawls, 2006). The true soil texture was defined as a silt loam (25% sand, 25% clay), which is among the common soil types for maize production in Nebraska (CropWatch, 2018). On each day of the simulation, water-flux inputs (rainfall, irrigation, and evapotranspiration) used to simulate farmer's expectations (i.e., Simulation 1) were also perturbed by normally distributed percentage errors (Figure 3.S4). The distributions chosen were unbiased (mean = 0) with standard deviations of increasing magnitude from 0 to 30% for both water-flux and soil-texture measurements. In the real world, measurement errors in these quantities may be on the lower end of this range, in particular in the context of quantities such as irrigation depths in more controlled application systems (e.g., centre-pivot or drip). However, we explore variations in errors up to 30% to provide an estimate of how measurement errors affect resulting irrigation water use and crop yields for a diversity of error levels ranging from perfect information (0% error) to very large uncertainty (30% error). Given that the added errors were randomly drawn, the process was repeated 1000 times for each combination of irrigation strategy, water-flux

error, and soil-texture error. This number was chosen by finding how many repetitions were required before average CE results stabilised (Figure 3.S5).

Simulations were conducted over a 30-year (1989–2018) period using weather data observed at a monitoring station in Champion, southwest Nebraska (HPRCC, 2016), where the number of years was chosen to capture the range of potential growing weather conditions experienced by farmers in the region. For each 30-year simulation, we determine the effects of soil-moisture uncertainty by comparing average annual profits and total irrigation water use for simulations with and without water-flux and soil-texture errors. Simulations consider a constant crop price used in Equation 3.2 of $180 per tonne ($4.57 per bushel) based on a 10-year average of US maize grain prices (USDA, 2019) and an irrigation cost of $2 per ha-mm ($20.55 per acre-inch) based on typical costs associated with pumping and application of water (UNL, 2017). Fixed production costs were set to $1728 based on typical average non-water production costs for centre pivot irrigated maize production in the 2019 Nebraska Crop Budget Report (CropWatch, 2019a). These costs include field operation costs including labour, materials and services, taxes, etc. for non-irrigation activities, and are not likely to change significantly if the irrigation strategy is altered. Baseline simulations consider a risk-neutral farmer (i.e., $r = 0$ in Equation 3.1).

To evaluate the sensitivity of our results to model parameter choices, we conducted a range of sensitivity tests. The first sensitivity analysis was for irrigation cost, whereby the same analysis was conducted for two alternative costs of irrigation ($1, $3 per ha-mm), as this cost can vary greatly due to differences in water source, energy price and tariffs, and labour costs (Wichelns, 2010). Higher irrigation costs will amplify the economic impacts of over-irrigation, potentially resulting in a corresponding increase in the value a farmer may place on accurate soil-moisture information. On the other hand, when irrigation is cheap to apply, there is less incentive to conserve water and risk any reductions in yield that may occur due to delayed irrigation. In these circumstances, the value of soil-moisture information therefore is expected to be lower.

Subsequently, we conducted a second sensitivity analysis to evaluate the effects of risk-aversion on the value of more accurate soil-moisture information to a farmer. Risk-averse farmers were represented by varying the risk coefficient in Equation 3.1 ($r = 0, 2.5, 5$) to reflect the fact that existing literature shows farmers often exhibit risk-averse behaviour

(Abdulkadri, 2003; Menapace et al., 2013). As more weight is placed on reducing the year-to-year variance in profits, the optimal strategy should aim to apply more water than the risk-neutral case, in order to reduce water stress in the driest years. The desire to reduce year-to-year variation should also increase the value the farmer places on accurate soil-moisture information, as greater information accuracy would be expected to result in lower risks of crop failure all else being equal. As well as isolating the effects of changing irrigation cost or risk preference on our chosen case study, the joint effects of changing both these parameters at the same time was also examined. The full sensitivity analysis is therefore made up of nine scenarios, representing three choices of irrigation cost and three choices of risk preference. A summary of key model parameters is shown in Table 3.2.

The experiments were performed on The University of Manchester's High Performance Computing (HPC) cluster. Up to 20 jobs could be ran in parallel with each job representing one full experiment (e.g., 1 strategy x 30 years x 8 flux errors x 8 texture errors x 1000 repetitions). Each job was computed on a 2×8-core Intel Xeon E5-2650 v2 @ 2.60GHz + 64GB RAM compute node. The time to complete all jobs for 50 optimized strategies, 3 irrigation costs and 3 risk preferences was approximately 1 week.

Table 3.1. Top five soil-moisture threshold strategies for baseline analysis (Risk coefficient *r* = 0; Irrigation cost *C* = $2 per ha-mm.).

THRESHOLD 1 (%TAW)	THRESHOLD 2 (%TAW)	THRESHOLD 3 (%TAW)	THRESHOLD 4 (%TAW)	CERTAINTY EQUIVALENT ($/HA)
48	46.3	31.8	11.6	191.9
44.8	46	33.7	0	191.8
44.9	48.4	33.6	5.9	191.5
44.1	49.3	32.7	0	189.7
42.8	47.4	33.8	8.7	189.1

Table 3.2. Overview of key model parameters.

PARAMETER	VALUE
Simulation years	1989–2018 (30 seasons)
Crop	Maize
Crop market price *M*	$180 per tonne ($4.57 per bushel)
Irrigation cost *C*	$ [1, 2, 3] per ha-mm ($ [10.28, 20.55, 30.83] per acre-inch)
Fixed costs *F*	$1728 per ha (CropWatch, 2019a)
Soil-texture error	0–30% standard error
Water-flux error	0–30% standard error
Optimization repetitions	50
Simulation repetitions	1000
Risk preference *r*	[0, 2.5, 5]
Maximum daily irrigation depth	25 mm

3.3 RESULTS

3.3.1 The effect of soil-moisture uncertainty on water use and profits

Each combination of water-flux and soil-texture errors (averaged over all strategies and repetitions), and their resulting impacts on water use and profits, is shown in Figure 3.3 (for a risk-neutral farmer and implementing behavioural feedbacks described in Section 3.2.2). The responsiveness of modelled outputs is variable depending on the underlying source of soil-moisture error, with water use (Figure 3.3a) and profits (Figure 3.3b) showing greater sensitivity to water-flux errors than soil-texture errors. Increasing magnitude of soil-moisture errors (from both sources) results in only marginal changes in irrigation water use and profits. For example, each contour line in Figures 3.3a and 3.3b represents a 1% increase and 1% decrease, respectively, from the case of perfect information. These results indicate that a nearly 30% assumed standard error in the water-flux and soil-texture inputs causes only around a 4% increase in total water use (Figure 3.3a) and 14% decline in profits (Figure 3.3b). Figure 3.S6 shows that almost no change in crop yields is observed, even for extreme levels of error of up to 30%, indicating that the primary driver of profit losses is through this excess water application through inefficient irrigation scheduling as a result of soil-moisture estimation uncertainty. The results presented in this section and following sections are averaged over the 50 CE-maximizing strategies, as well as 1000 repetitions of the 30-year case study simulation. Thus, the impact of the stochastic measurement uncertainty may deviate from the overall trend in any given simulation (Figure 3.S7).

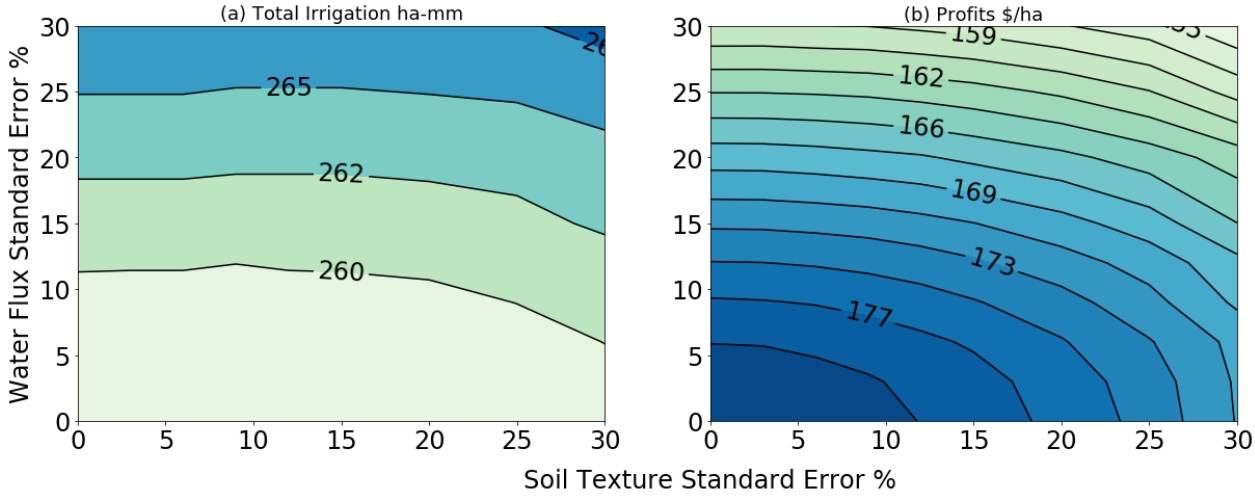

Figure 3.3. The effect of increasing water-flux and soil-texture errors on (a) total irrigation [ha-mm], and (b) profits [$ ha^{-1}]. The percentage errors are drawn from a normal distribution of mean = 0 and standard deviation increasing from 0–30%. The contour lines represent a 1% change from the case of zero added error. The effect of added uncertainty in water-flux and soil-texture measurements is almost no change in yields despite marginal increases in irrigation. This contrast implies irrigation is applied sub-optimally, leading to profit losses from wasted water application. (Risk coefficient r = 0; Irrigation cost C = $2 per ha-mm.)

Figure 3.4 also illustrates the effect of adding the additional behavioural-feedback mechanisms described in Section 3.2.3. Behavioural feedbacks were designed to prevent unrealistic irrigation behaviour caused by measurement error, effectively ensuring that irrigation events were triggered when true soil-moisture levels were sufficiently low to cause visible plant stress and were not triggered when true soil-moisture levels were sufficiently high to cause visible soil saturation. Introducing the feedback mechanism leads to greater increases in water use as soil-moisture errors become larger (Figure 3.4a), but results in a reduction in the economic losses caused by larger measurement errors (Figure 3.4b). This result reflects the fact that the behavioural feedbacks most commonly act to trigger irrigation when the farmer has underestimated soil-moisture depletion, leading to more frequent irrigation events but also reducing the duration and severity of crop water stress relative to simulations without behavioural feedbacks. In contrast, triggering of behavioural feedbacks for excess irrigation is relatively uncommon as the whole root-zone would have to have exceeded field capacity.

Introducing behavioural feedbacks also results in small changes in water use and reductions in CE and profits even under perfect information (i.e., for errors equal to 0). Because CE-maximizing strategies are designed under the assumption of perfect information without consideration of additional behavioural feedbacks, introducing these feedbacks leads to permutations to these CE-maximizing strategies that trigger slight adjustments to final simulated outputs at the end of each season. These changes, however, are small, representing as little as 0.5% and 0.3% changes in average annual water use and CE, respectively.

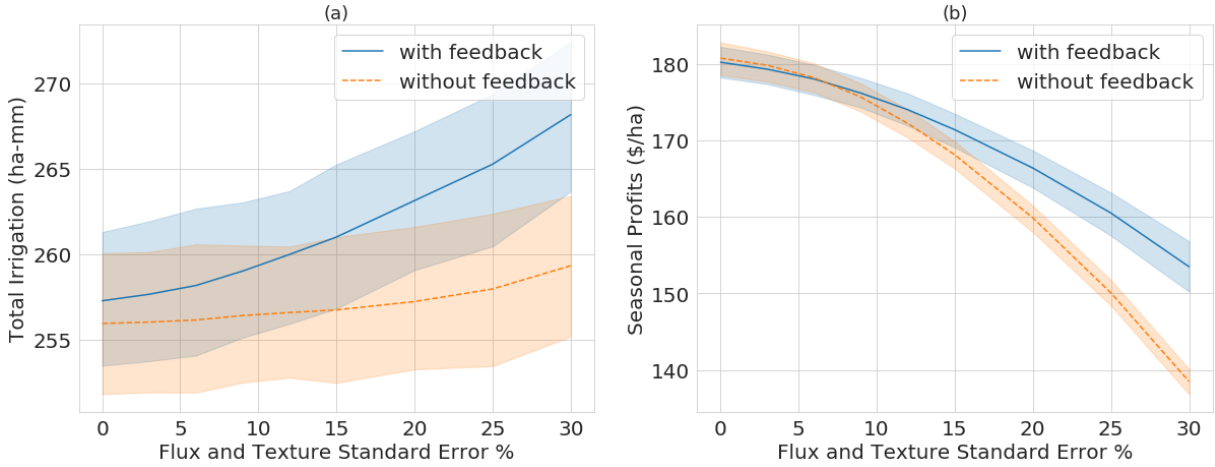

Figure 3.4. The effect of increasing water-flux and soil-texture error on (a) total irrigation [ha-mm] and (b) profits [$ ha^{-1}] when the magnitude of both sources are equal (i.e., plotting the bottom-left to top-right diagonal of Figures 3.3a, 3.3b). The shaded area represents the 95% confidence interval over all 50 strategies. The comparison with and without the feedback mechanism is also presented to show the effect that simple farmer behaviour adjustments can have on the value of soil-moisture information. (Risk coefficient r = 0; Irrigation cost C = $2 per ha-mm.)

3.3.2 Sensitivity analysis

To assess the impact of our choice of risk preference and irrigation cost (r = 0 and C = $2 per ha-mm) on the estimated impacts of soil-moisture uncertainty on irrigation water use and farm profits, we repeat the analyses described in Sections 3.2.2 and 3.2.3 for combinations of three alternative risk coefficients (0, 2.5, 5) and irrigation costs ($1, 2, 3 per ha-mm). In the absence of soil-moisture uncertainty, higher irrigation costs lead to reductions in average water use as expected (Figure 3.5a). The magnitude of this reduction is also much larger than increases resulting from risk preferences or soil-moisture uncertainty (Figure 3.5b), suggesting that irrigation-water costs or prices have a larger effect on water-use decisions than risk aversion or uncertainties in underlying data used to determine soil-moisture conditions for irrigation scheduling.

This relationship between irrigation cost and water use is commonly analysed in the literature via the *price elasticity of water use* (also known as the price elasticity of demand), which is the percentage change in a farmer's irrigation water use for a percentage change in irrigation cost. Previous studies have placed the *price elasticity* within a wide range 0.001–1.97 (Scheierling et al., 2006). Calculating the price elasticity of water use using our results in the absence of soil-moisture uncertainty yields a mean *price elasticity* of 0.23 ± 0.12 (calculated from the results presented in Figure 3.5a). This result is comparable with past estimates that consider only water-use adjustments on the intensive margin (i.e., estimates that consider changes to per-area irrigation water use, but not changes in demand caused by adjustments to irrigated area or crop type). Similarly, calculating the percentage change in water use for a percentage change in soil-moisture error or risk coefficient yielded an *information elasticity* of 0.015 ± 0.016 and a *risk elasticity* of 0.012 ± 0.016. This comparison further reinforces our finding that water use appears to be much more sensitive to changes in cost than either risk preference or soil-moisture error.

The choice of irrigation cost and risk preference can also alter how soil-moisture uncertainty affects water use and profits in the case of farmers with an imperfect ability to estimate or monitor soil-moisture conditions. High irrigation costs will reduce the optimal amount of irrigation used over the season, meaning that farmers have less scope to mistime events while still avoiding damaging water stress or that farmers must apply extra irrigation at

significant economic cost. A high degree of risk aversion (represented by a larger positive value of *r*) will similarly penalise large year-to-year fluctuations in profit, meaning the negative economic impacts of soil-moisture errors (and hence the value of improved information accuracy) may be enhanced due to the increase in the variance in annual profits. When both high risk preferences (r = 5) and high irrigation costs (C = \$3 per ha-mm) are combined, the economic losses resulting from uncertainty are almost doubled compared to the case presented throughout this analysis (r = 0 and C = \$2 per ha-mm) (Figure 3.5c).

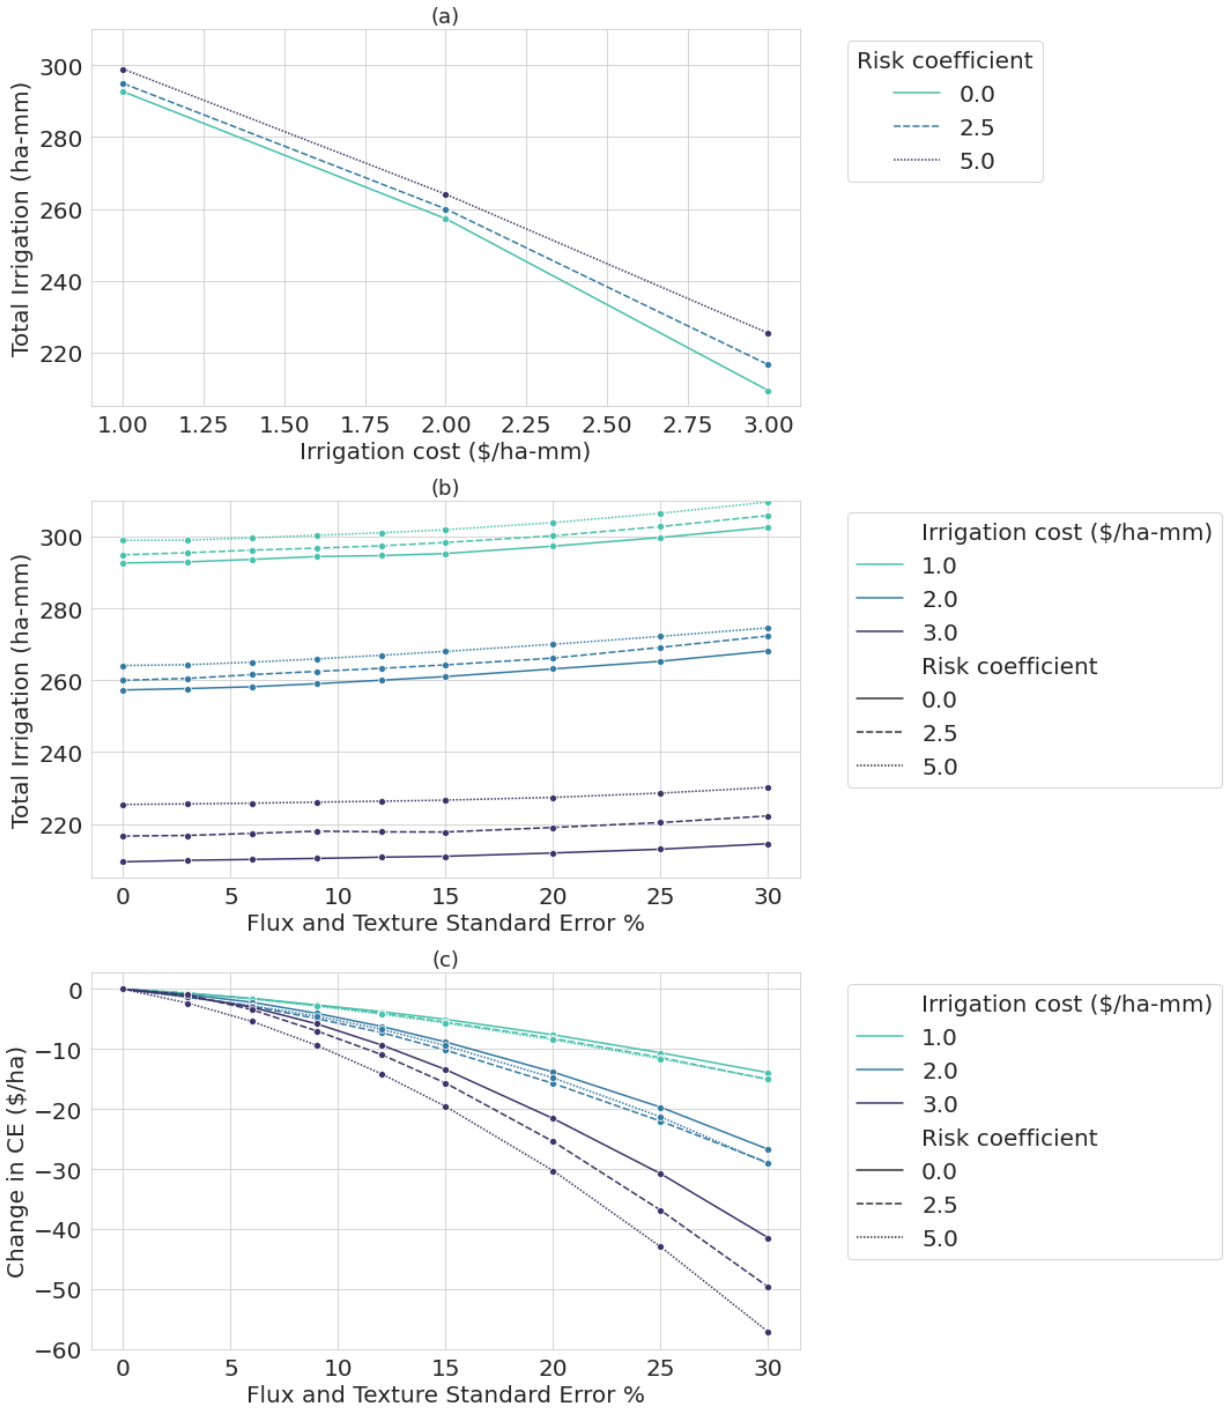

Figure 3.5. The effect of increasing irrigation costs [$/ha-mm] on water use in the absence of water-flux and soil-texture error (a). The effect of increasing water-flux and soil-texture error on (b) seasonal irrigation and (c) certainty equivalent for varying irrigation costs and risk coefficients when the magnitude of both sources are equal.

3.3.3 Comparing the impacts of uncertainty with choice of irrigation strategy

Our analysis in the previous sections has shown that a 30% water-flux and soil-texture error results in only a 4% increase in water use (Figure 3.3a), whereas the standard deviation in water use between 50 different 'optimal' strategies was 5.7% (Figure 3.4a). This comparison indicates that the choice of irrigation management strategy (represented in this study by a combination of four soil-moisture thresholds) could be more important in determining irrigation-water use than uncertainty in underlying soil-moisture estimates used to implement irrigation strategies. Since all strategies evaluated so far can be considered 'optimal', one would expect a 'sub-optimal' strategy to have a significantly different level of water use and profits if that were the case.

To explore further the impacts of the choice of irrigation strategy relative to the effects of soil-moisture uncertainty, we repeated the analysis for our case study, allowing the farmer to select only a single optimised soil-moisture threshold throughout each season (e.g., the green line in Figure 3.2 would be one straight line) as opposed to a set of four optimised thresholds disaggregated by growth stage. The top five single-threshold strategies in terms of 30-year average CE are displayed in Table 3.3. Single thresholds have been applied in past studies of farmer irrigation decision-making and is often recommended by extension and agronomic advisories (Blonquist et al., 2006; Gutierrez et al., 2014). However, since they do not consider the variable sensitivity of crop growth to water deficits over the growing season, they are sub-optimal compared to the four-threshold strategies used in the baseline analysis.

Figure 3.6 illustrates that the choice of irrigation management strategy (i.e., one versus four soil-moisture thresholds per season) has a much larger impact on water use (Figure 3.6a) and profits (Figure 3.6b) than added soil-moisture uncertainty. Under perfect information (zero added soil-moisture error), the sub-optimal strategy (one irrigation threshold) uses 35 ha-mm more water on average and reduces profits by \$117 ha^{-1} compared with the optimal strategy (four irrigation thresholds). Conversely, adding 30% water-flux and soil-texture errors to the optimal strategy results in only 11 ha-mm more water being applied and a \$27 ha^{-1} reduction in profits. This means that the accuracy of soil-moisture information appears to be much less important to the water-use efficiency and profitability of irrigation decision-making than the choices of heuristics for irrigation scheduling. Indeed, Figure 3.6 suggests

that adding a 30% standard error onto an optimal strategy can still outperform a sub-optimal strategy with perfect information.

Table 3.3. Top five soil-moisture threshold strategies for a single seasonal threshold (Risk coefficient *r* = 0; Irrigation cost *C* = $2 per ha-mm.).

THRESHOLD 1 (%TAW)	THRESHOLD 2 (%TAW)	THRESHOLD 3 (%TAW)	THRESHOLD 4 (%TAW)	CERTAINTY EQUIVALENT ($/HA)
38.2	38.2	38.2	38.2	65.2
35.9	35.9	35.9	35.9	64.2
35.1	35.1	35.1	35.1	64.1
35.1	35.1	35.1	35.1	63.7
40.1	40.1	40.1	40.1	63.4

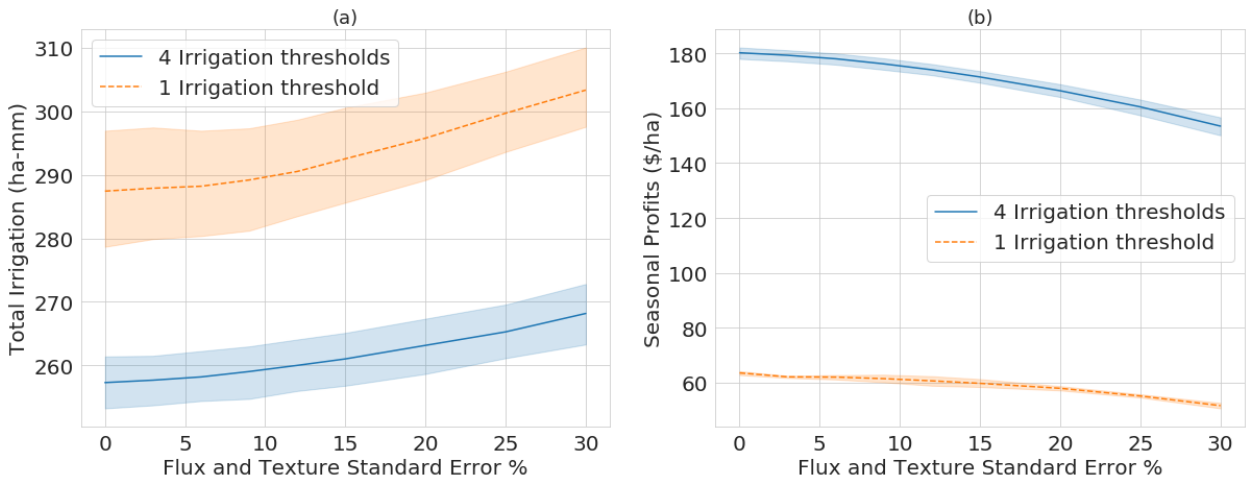

Figure 3.6. Comparison of (a) water use and (b) profits under increased uncertainty for two different types of irrigation strategy. The first strategy consists of four irrigation thresholds corresponding to the allowable level of soil moisture during four major growth stages. The second strategy is for one constant irrigation threshold that does not vary over the season. Particle-swarm optimization found both strategies to be CE maximizing (Risk coefficient r = 0; Irrigation cost C = $2 per ha-mm.)

3.4 Discussion

Our findings demonstrate that the response of irrigation-water use and farm profits to soil-moisture uncertainty is non-linear, with soil-moisture uncertainty leading to reductions in both efficiency of irrigation water use and farm profits. However, we show that, with the addition of realistic farmer behaviour (in the form of our feedback mechanism), the magnitude of these changes is much smaller than commonly assumed in policy and practice (Soulis et al., 2015; Adeyemi et al., 2017). *In fact, our analysis suggests that near-optimal irrigation decisions can be made without perfect soil-moisture information.* In contrast, we find the choice of irrigation management strategy – in this study given by a set of four soil-moisture thresholds – was a far more important determinant of water efficiency, productivity, and profitability than the accuracy of soil-moisture information used to implement a selected irrigation management strategy.

To explain this counterintuitive result, it is necessary to consider how the choice of irrigation management strategy and the accuracy of the information used to implement that strategy, combine to influence risks of crop water stress and resulting yield losses. For our baseline analysis, the highest performing (in terms of CE) soil-moisture target strategy consisted of irrigation being triggered when soil-moisture levels were assumed to be 48, 46, 38 and 12% of soil-water holding capacity in each of the crop's four main growth stages (Table 3.1). These assumptions compare with calibrated thresholds for the initiation of crop water stress affecting canopy expansion (86%), plant transpiration (31%), canopy senescence (31%), or pollination (20%). Thus, the optimal soil-moisture targets, apart from that in the final growth stage (canopy senescence) are above thresholds needed to trigger more severe crop stress effects (stomatal closure, early canopy senescence, pollination failure). These thresholds have been optimised to maximise profit over 30 years of climate data, meaning that they must incorporate a degree of headroom in order to cope with unexpected year-to-year weather variability, with exact weather conditions being unknown to farmers apriori in any given season. Therefore, errors in estimates of soil-moisture conditions have to be large and persist for multiple days to result in unintended water stress, reducing the risks posed by inaccurate soil-water monitoring. In contrast, the highest performing single threshold (i.e., sub-optimal) irrigation strategy was 38% of soil-water holding capacity (Table 3.3). This sub-optimal threshold is closer to the crop water stress thresholds – in particular in earlier

stages of crop development – and so can result in greater increases in both the frequency and magnitude of water-stress occurrence. This single seasonal threshold also has the effect of triggering excess water use in the late season, explaining how the choice of irrigation strategy has a large impact on water-use efficiency and farm profits even in the absence of soil-moisture uncertainty.

This intuition underlying our findings is also critical to understanding how consideration of behavioural feedbacks (Section 3.3.1) and variability in irrigation costs (Section 3.3.2) influence the impacts of soil-moisture uncertainty on water use and farm profits in our analysis. When considering the potential for rational behavioural feedbacks made by farmers (e.g., to respond to visible signs of severe water deficits), it is logical that this feedback mitigates the effect of measurement uncertainties as these feedbacks further minimise any risks of erroneously falling below the intended irrigation trigger threshold. Similarly, the moderate increase in the sensitivity of water use and profits to soil moisture-uncertainty at higher water prices can be explained by the fact that farmers faced by higher water prices will irrigate at lower soil-moisture levels independent of any uncertainty in soil-moisture measurement (in order to reduce total irrigation applied over the season). Indeed, optimal soil-moisture targets for higher water price of $3 per ha-mm were 0, 46, 30, 2% soil-water holding capacity. As such, farmers operating in regions with higher water prices – whatever their cause – will have greater risks of triggering unintended water stress for a given level of soil-moisture measurement, suggesting that the value of improved soil-moisture information accuracy will depend in part on local water prices and costs of irrigation access.

3.4.1 Policy Implications

Our main finding that perfect soil-moisture information is not required to make near-optimal irrigation decisions is broadly consistent with wider literature assessing the value of soil and climate information for farmers (Bosch & Eidman, 1987; Botes et al., 1996; Fafchamps & Minten, 2012). For example, previous research found that using historical weather data to determine optimal irrigation strategies results in relatively minor irrigation increases and profit losses compared with using perfect weather forecasts (Linker et al., 2018; Jamal et al., 2019). Moreover, Linker & Kisekka (2017) used a similar parallel-model

approach to the present study to show that perfect real-time soil-moisture measurements may not be required to implement a deficit irrigation strategy.

These studies, along with our findings, show that the value of improved soil-moisture measurement data alone may have been overstated by technology providers and water managers. Instead, a meaningful proportion of potential gains in profit or water efficiency may actually be the result of improved advice or guidance about optimal irrigation management practices, which are commonly provided by technology providers or extension agents when seeking to introduce new technologies for farm-level decision support (whether that be a soil-moisture sensor or crop-water model). When compared with estimated profitability gains observed in our analysis, this result suggests that the benefits of improved soil-moisture data – which can cost hundreds of dollars (Kukal et al., 2019) – may vary substantially depending on farmers' baseline soil-moisture uncertainty or irrigation scheduling practices and heuristics, with much of the gains in water-use productivity potentially achievable through improving irrigation management practices, even if soil-moisture monitoring remains imperfect.

The greater impact of irrigation-strategy choice over information accuracy identified in our study has important implications for efforts to improve irrigation-management decision-making. For example, deficit irrigation – whereby irrigation is reduced below full crop-water needs while minimizing stress at sensitive growth stages – has been widely proposed as a key tool for improving agricultural water productivity (Kang et al., 2000; Blonquist et al., 2006; Fereres & Soriano, 2006; Dukes et al., 2010). Our findings highlight the importance of these efforts, demonstrating that it may be more cost effective to refine irrigation-management strategies, for example through optimization methods, in particular where costs of reducing soil-moisture monitoring errors are large (e.g., if multiple sensors need to be installed within a single field). A focus on improving irrigation strategies over monitoring data could lead to potentially faster improvements in agricultural water productivity. Where improvements in monitoring are considered, the focus should be on low cost and more scalable solutions as opposed to expensive sensor network arrays that aim for high levels of accuracy and precision in soil-moisture monitoring. For example, a farmer might choose to use a smaller network of low-cost sensors to monitor soil moisture or instead rely on remotely observed proxies for soil moisture such as from Earth observation satellites. Our

results, and those of Linker & Kisekka (2017), indicate that the lower monitoring accuracy of such approaches would not necessarily be detrimental to water-use efficiency or productivity, and in some cases may be an economically optimal solution for farmers. Further research, however, is required to evaluate the cost-effectiveness of different monitoring and sensing approaches that vary in terms of their adoption costs and accuracy in different types of agricultural production systems.

Finally, our results contribute to the wider literature on the use of crop-water models to estimate irrigation water demand in the context of agricultural water management and planning. Such studies typically assume the farmer has access to perfect soil-moisture information (García-Vila & Fereres, 2012; Foster et al., 2014; Zellner et al., 2020). Our results suggest that this assumption does not have a major impact on results other than marginal underestimations in water use. However, our results do highlight that incorrectly parameterizing farmers' irrigation strategies is potentially a significant source of uncertainty in irrigation demand projections. Previous studies have highlighted a large variability in irrigation behaviour, even after controlling for key inputs to crop models such as crop type, soil type, and irrigation technology (Foster et al., 2019). Our findings thus highlight that the lack of information on the rules that govern farmers' irrigation decisions is potentially a major source of uncertainty in irrigation water demand estimation and projection. This uncertainty may greatly exceed those introduced by uncertainty in soil, climate, or crop input parameter data that are more commonly the focus of parameter uncertainty studies.

3.4.2 Limitations and Future Work

The analysis presented in this study contains a number of limitations and simplifications that warrant discussion here and consideration in future extensions to this paper. In our analysis, we have assumed a perfect model that accurately represents both the movement of water in soil and corresponding crop growth. This assumption has enabled us to focus on quantifying the effect of measurement errors, such as the water flux and soil texture, making the results independent of the choice of crop-growth model. If this assumption does not hold, there would be a larger deviation between the farmer's perception and the true soil-water conditions, potentially placing more importance on having accurate soil-moisture information. However, the inclusion of the feedback mechanism in this study places limits on how wrong these perceptions of soil-water conditions can be, meaning that the main

conclusions of this study should hold even under an imperfect model. Moreover, in reality, adoption of sensors is unlikely to guarantee perfect information about soil-moisture conditions. Sensors may give imprecise readings due to faults, inaccurate calibration, or where extrapolating point measurements to field scales (Sharma et al., 2021), meaning that our comparison with perfect information may be exaggerating the true benefits of technology adoption.

The major consequences of over-irrigating in our framework have been through the excess cost of water applied. However, in reality, an additional impact of over-irrigation is waterlogging and associated damage to crops through aeration stress, pests/diseases and mechanical damage such as lodging. Recent research has shown that process-based crop models poorly capture the occurrence of waterlogging due to assumptions in soil-moisture drainage routines, and thus underestimate the negative impacts of excess water supply on crop growth and yields (Li et al., 2019). Neglecting these negative consequences of over-irrigation may mean that our analysis has understated the effects of uncertainty on crop yields and profits, in particular for crops that are sensitive to effects of waterlogging and excess irrigation (e.g., enhanced lodging risks for crops such as wheat). Nonetheless, the more severe cases of over-irrigation would have been prevented via the behavioural feedback mechanisms used in this analysis. and This point of intervention, where a farmer feels that further irrigation would be detrimental to the crop (chosen to be field capacity in this case study), would likely change and reflect the risks of their specific crop based on their own experience. This feedback addition means that impacts of model simplifications of over-irrigation effects should be minimised under the assumption of rational farmer behaviour. Conversely, for systems that apply significant amounts of water in one event (e.g. flood irrigation), the potential for crop stress caused by over-irrigation is increased compared to our study of pivot production.

Our analysis does not explicitly consider the effects technical or regulatory constraints to irrigation water use on the value of more accurate soil-moisture information. In reality, policies constraining seasonal or intraseasonal rates of water use (Ifft et al., 2018) or physical limits on water use (e.g., due to declining well yields) (Foster et al., 2015) may both increase incentives for farmers to avoid over-irrigation and thus increase the value of implementing optimal irrigation strategies. Combined with these limits, characteristics of

the delivery system will impact which days and how much water can be applied in each event beyond the simple maximum daily limit considered in our analysis (O'Brien et al., 1998). Such restrictions and extraneous variables (e.g., a seasonal cap on irrigation) may increase the value of information as there is an additional cost of exceeding this cap. On the other hand, if irrigation events are largely determined by well yields, water delivery systems, energy regulators, and weather forecasts, then the number of mistimed irrigation events due solely to inaccurate soil-moisture measurements will be low. In either case, we expect the value of improved irrigation strategies to move in the same direction as the value of soil-moisture information, as more freedom over the timing and quantity of water applied allows the farmer to better adjust irrigation to match crop-water needs.

The results of our illustrative example are predicated on the ability of the underlying crop-water model (AquaCrop-OS) to accurately simulate crop responses to soil water deficits. The model has been calibrated and extensively applied to simulate irrigated and rainfed maize production in our study area and wider North American farming systems (Heng et al., 2009; Foster et al., 2015; Sandhu & Irmak, 2019). However, to ensure that our results are not being driven by any remaining uncertainty in either water-stress thresholds or response parameters, a sensitivity analysis was therefore performed. In this sensitivity analysis, we artificially altered the stress-response parameters in AquaCrop-OS, setting each value to 1 to maximise impacts of any short-duration periods of water stress on crop growth and yields. We then repeated our analysis described in Section 3.3.1 following methods outlined in Section 3.2. Water use and profits from this analysis showed the same non-linear response to water flux and soil-texture uncertainty as reported earlier in Section 3.3. A 10% standard error in both soil-texture and water-flux inputs led to less than 1% change in water use and less than 5% change in profits (Figure 3.S8). This sensitivity analysis demonstrates that the results presented in this article are highly unlikely to be an artefact of any inability within AquaCrop to simulate the severity of short-duration water-stress impacts, and we expect that qualitatively identical results would therefore be obtained with alternative more biophysically complex crop models [e.g., Hybrid Maize (Yang et al., 2004)]. A key reason for this result is that while the model now responds more strongly to water stress, optimal soil-moisture targets are also adjusted to 64, 73, 39, 0% soil-water holding capacity and thereby greatly increasing the headroom above the critical crop-stress thresholds.

Finally, our results are specific to the particular illustrative application chosen in this study: irrigated maize production in the central United States. Nonetheless, we expect that our general qualitative finding – that choice of irrigation scheduling rules is more important than the accuracy of data these rules are implemented based on – will apply more generally to other crops and regions. In particular, our sensitivity analysis (described above) suggests that, even for more drought-sensitive crops (e.g., soft fruits), perfect information may not be necessary to achieve near-optimal water efficiency and profits so long as farmers' baseline irrigation heuristics are effectively representative of crop sensitivity to water deficits. Conversely, higher crop prices are likely to enhance the value of improved soil-moisture information as any yield losses will have a larger economic impact for farmers. The limited sensitivity of crop yields to measurement errors shown in Sections 3.3.1 and 3.3.2 suggest that changes in crop price would have minimal impact on economic losses caused by soil-moisture uncertainty. This is in contrast with the larger effects of water prices and risk aversion that are directly affected by the greater magnitude of change in irrigation water use with measurement error. Nonetheless, we acknowledge that variations in crop yields may have a larger influence on the value of improved soil-moisture information, for example where a farmer's baseline irrigation strategy is sub-optimal or for crops that respond strongly to either soil water deficits or saturation. Analysing such alternate production environments is beyond the scope of this paper, but will be an important area for future research building on the results presented here.

3.5 Conclusions

Effective irrigation scheduling is essential for improving agricultural water productivity and managing the negative impacts of agricultural water abstractions on other water users and the environment. Therefore, technology companies and researchers have focused on developing technologies such as soil-moisture sensors that improve the quality of farmers' soil-moisture information. Implicit to this approach is an assumption that more accurate information is essential to enable efficient near-optimal irrigation decisions. However, this assumption neglects the contribution of farmers' choice of irrigation strategy, which may have an equally significant or greater impact on overall efficiency and profitability of agricultural water use.

In this article, we have developed a framework to assess the impacts of increasing uncertainty on water use and farm profits. Measurement errors in water fluxes and soil texture lead to a divergence between the farmers' perception of soil moisture and the true soil-water status, resulting in sub-optimal irrigation decisions and reduced profitability. However, our results show that the magnitude of these impacts are small, with a 30% standard error in water-flux and soil-texture measurements – which is much larger than the likely real-world measurement uncertainty in these quantities, particularly precipitation and irrigation depth – resulting in only a 4% increase in average water use, meaning that near-optimal irrigation decisions can be made without perfect information. Contrastingly, we demonstrate that the choice of irrigation scheduling strategy has a larger impact on water use and profits than soil-moisture uncertainty.

Our analysis suggests that efforts to improve irrigation water efficiency should therefore focus primarily on helping farmers to evaluate and develop improved irrigation scheduling strategies for their specific production settings. Where existing irrigation management strategies are poorly aligned with local agronomic and biophysical conditions, the potential gains in water-use productivity and profitability may be large and further enhance the potential benefits from complementary efforts to improve quality of soil-moisture information using new forms of low-cost and scalable sensing technologies.

Acknowledgments, Samples, and Data

We thank the anonymous reviewers for comments that have improved the manuscript.

The work contained in this article was funded by the National Environmental Research Council's Understanding the Earth, Atmosphere, and Ocean Doctoral Training Programme, Grant NE/L002469/1.

Data generated by the developed framework and code can be found at http://doi.org/10.5281/zenodo.4041476.

SUPPORTING INFORMATION

This supporting information includes Figures 3.S1 through 3.S7. Figure 3.S1 illustrates how the feedback mechanism is triggered in an example season. Figures 3.S2 and 3.S5 provide insight into the choice of total optimizations and repetitions. Figures 3.S3 and 3.S4 illustrate how the added errors impact the soil-texture and water-flux measurements. Figure 3.S6 shows the effect of measurement uncertainty on yields. Figure 3.S7 shows the sensitivity of various irrigation strategies to the addition of measurement uncertainty. Figure 3.S8 shows the effect of measurement uncertainty on water use and profits for a highly drought-sensitive maize crop.

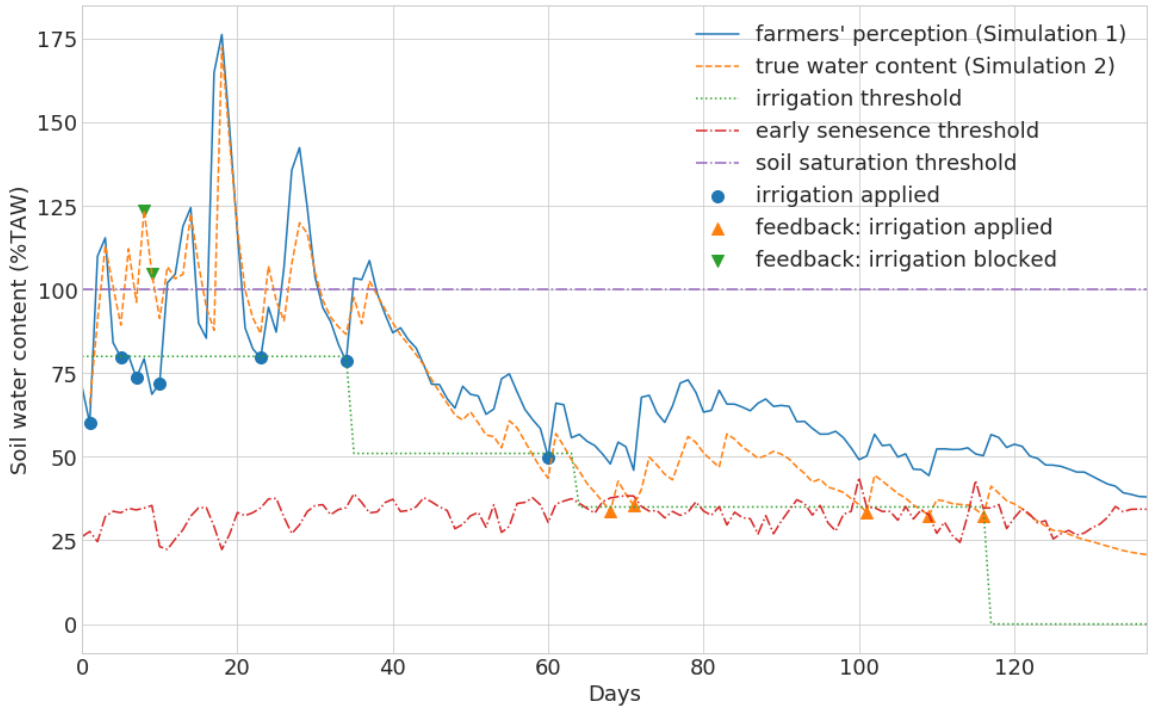

Figure 3.S1. Soil-moisture content (% total available water) during an example simulation, incorporating feedback elements to farmer irrigation decisions. Each day, irrigation is triggered (blue dots) when the farmers' perception (blue solid line) drops below the pre-defined threshold (green dotted line). The system feedback prevents these irrigation events occurring if the previous days' true water content (orange dashed line) was above field capacity (purple dot-dashed line), as it is assumed a farmer would be able to see if the soil were visibly wet, and hence choose not to irrigate (green downward triangles). The feedback mechanism also triggers irrigation if the true water content is low enough to trigger early senescence of the crop canopy (orange upward triangles). This trigger is based on the assumption that the crop would be showing visible signs of water stress (e.g., wilting of leaves), causing the farmer to irrigate unless they seek to allow water stress at this stage of the season (e.g., due to limited total water supply).

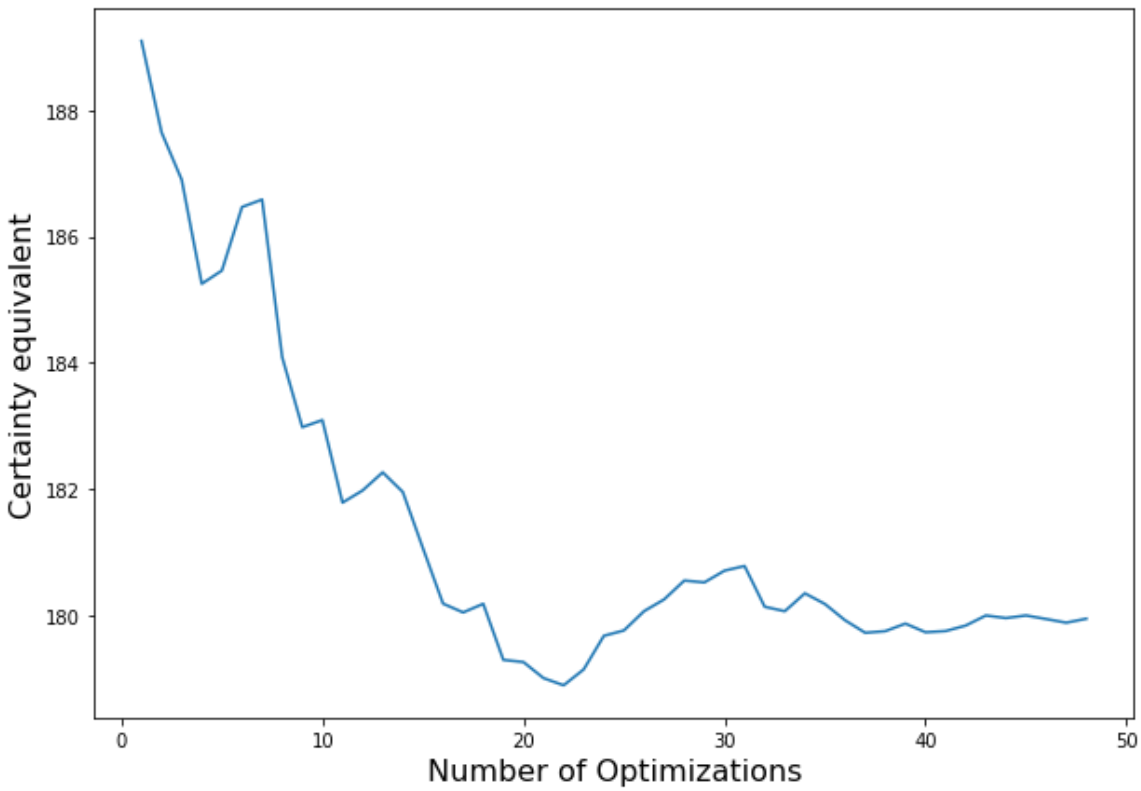

Figure 3.S2. The certainty equivalent in the absence of soil-moisture uncertainty averaged over all optimizations.

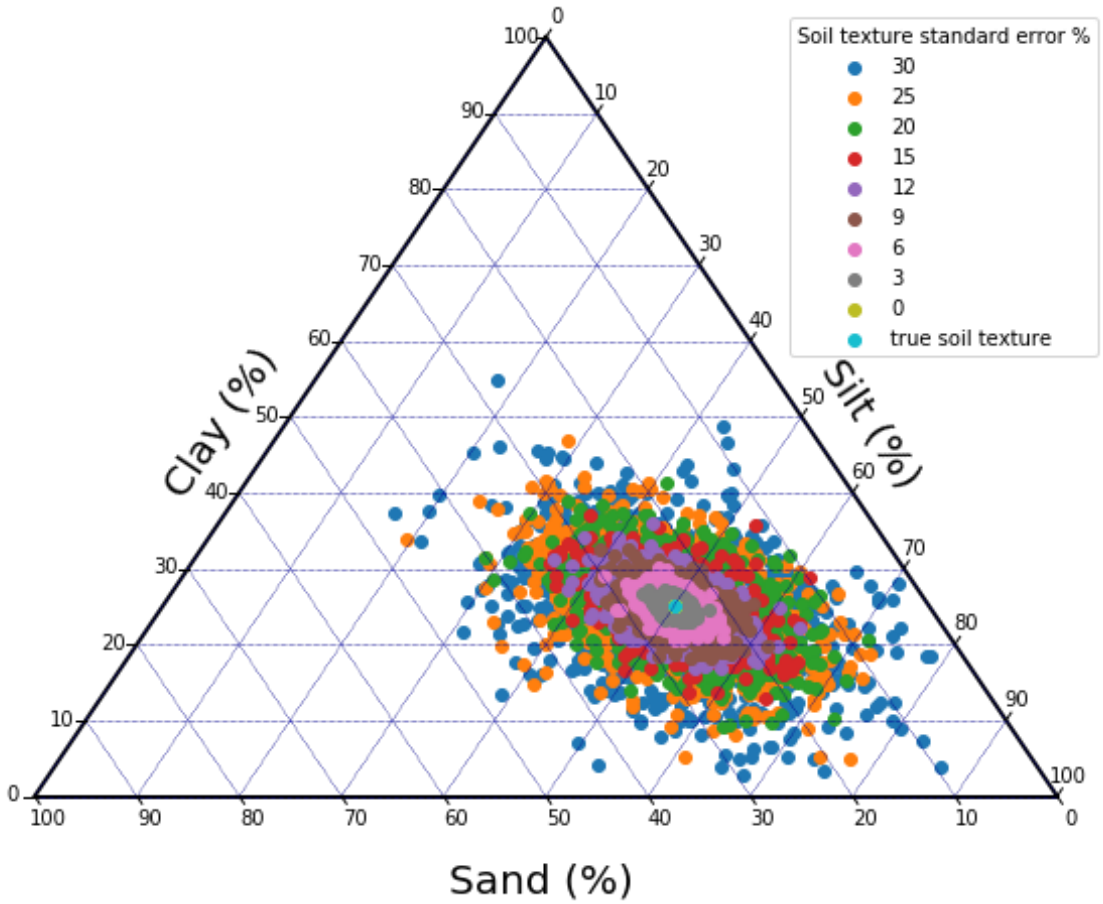

Figure 3.S3. The soil-texture composition after adding various magnitudes of percentage errors. The true soil composition is 25% sand, 25% clay, 47.5% silt and 2.5% organic matter (this figure ignores the organic matter content and combines it with silt content).

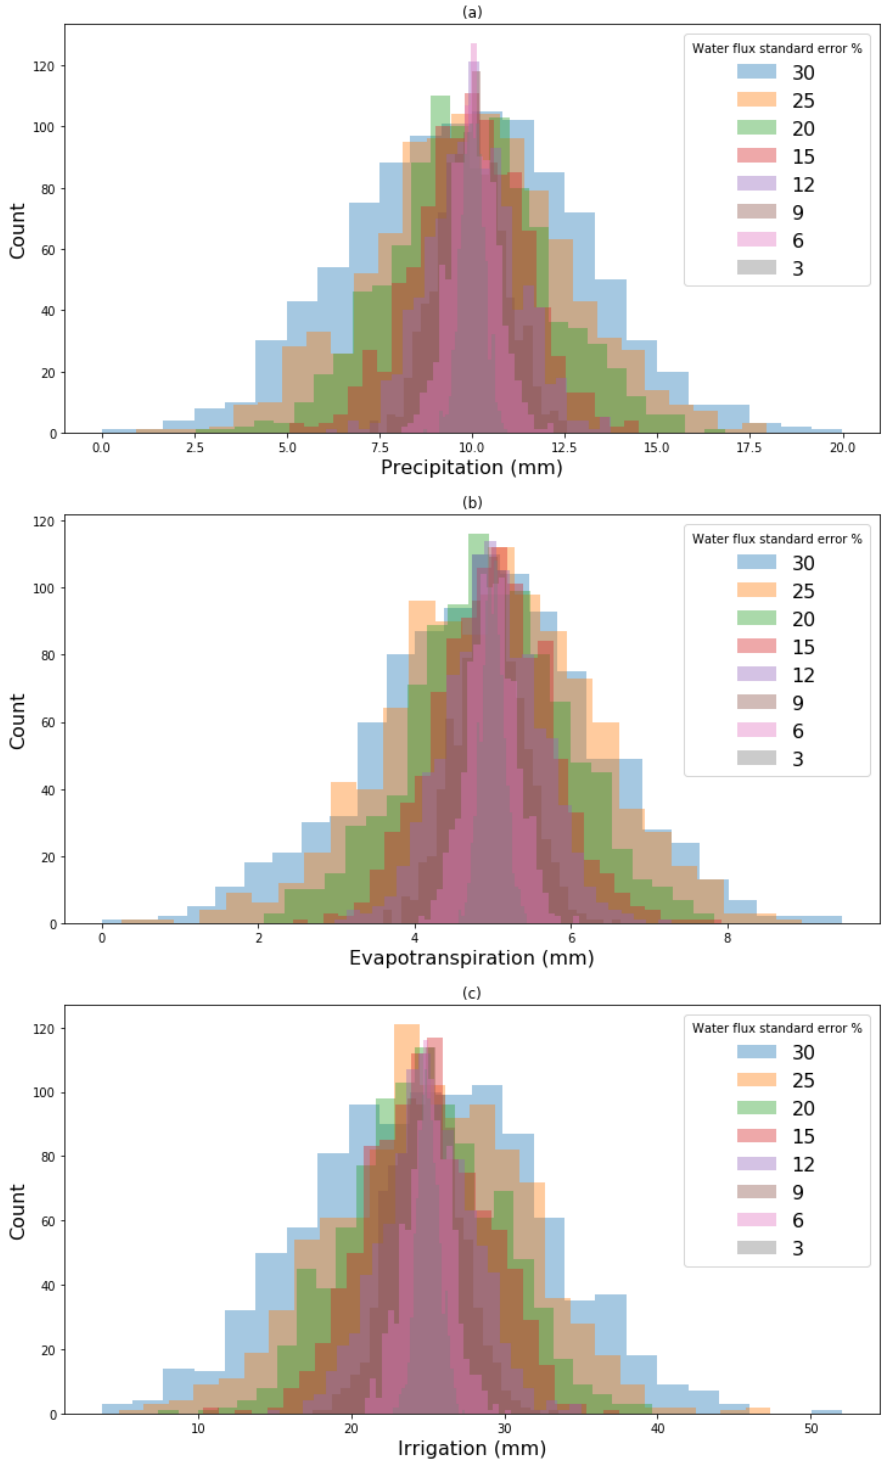

Figure 3.S4. The estimated (a) daily precipitation, (b) evapotranspiration, (c) irrigation depth after adding various magnitudes of percentage errors. For this example day, the true precipitation, evapotranspiration and irrigation applied were 10 mm, 5 mm and 25 mm respectively. The distributions are over 1000 repetitions for each magnitude of error.

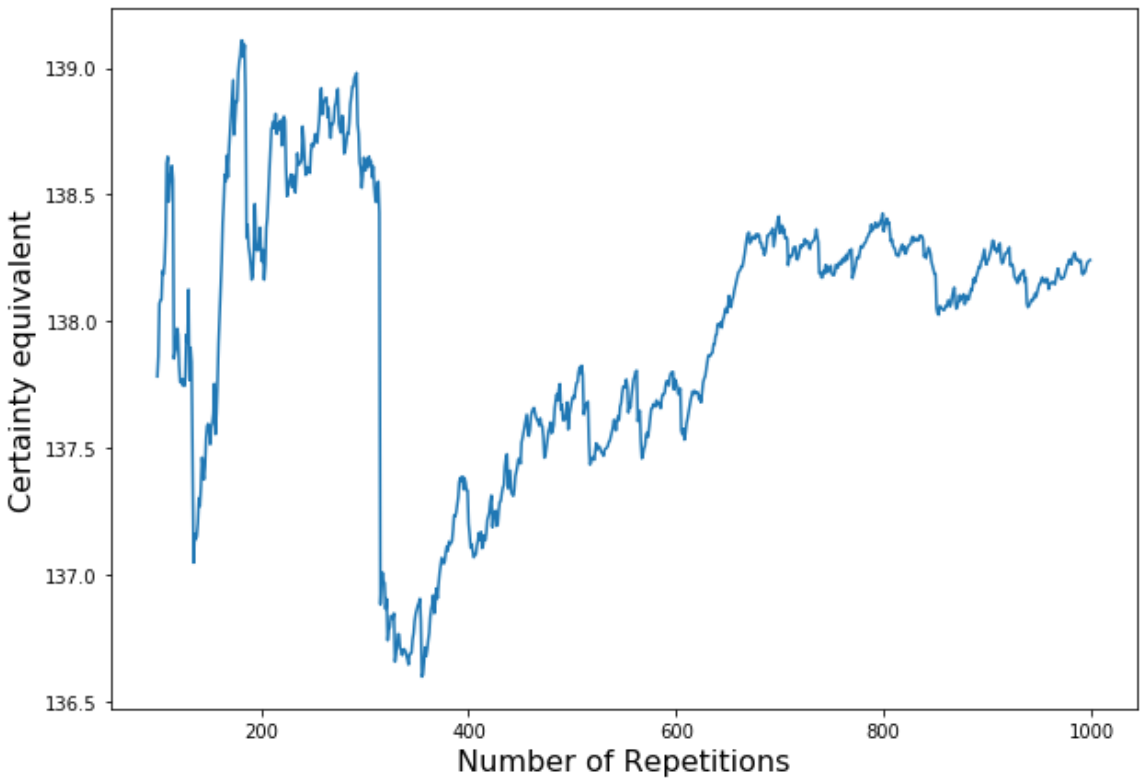

Figure 3.S5. The certainty equivalent for a given irrigation strategy (set of thresholds) averaged over all repetitions.

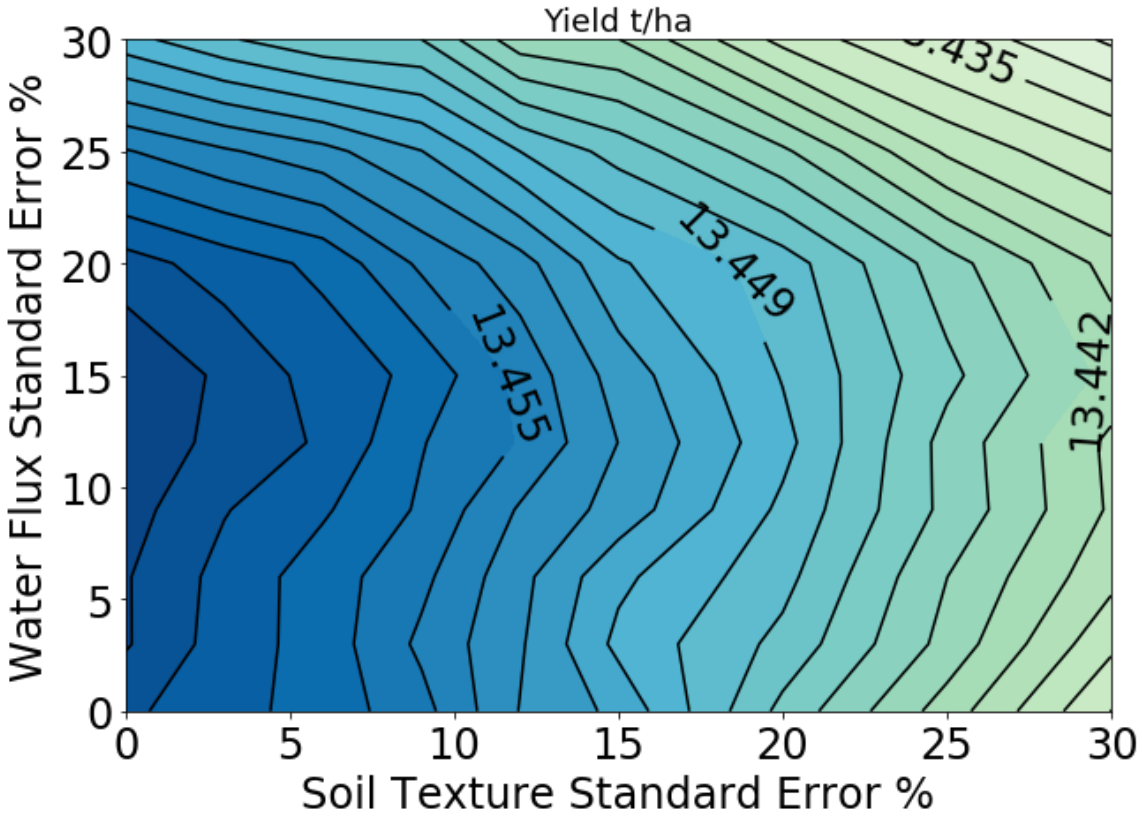

Figure 3.S6. The effect of different magnitudes of water flux and soil texture error on average yields. Each contour line represents a 0.01% change from the case of zero errors.

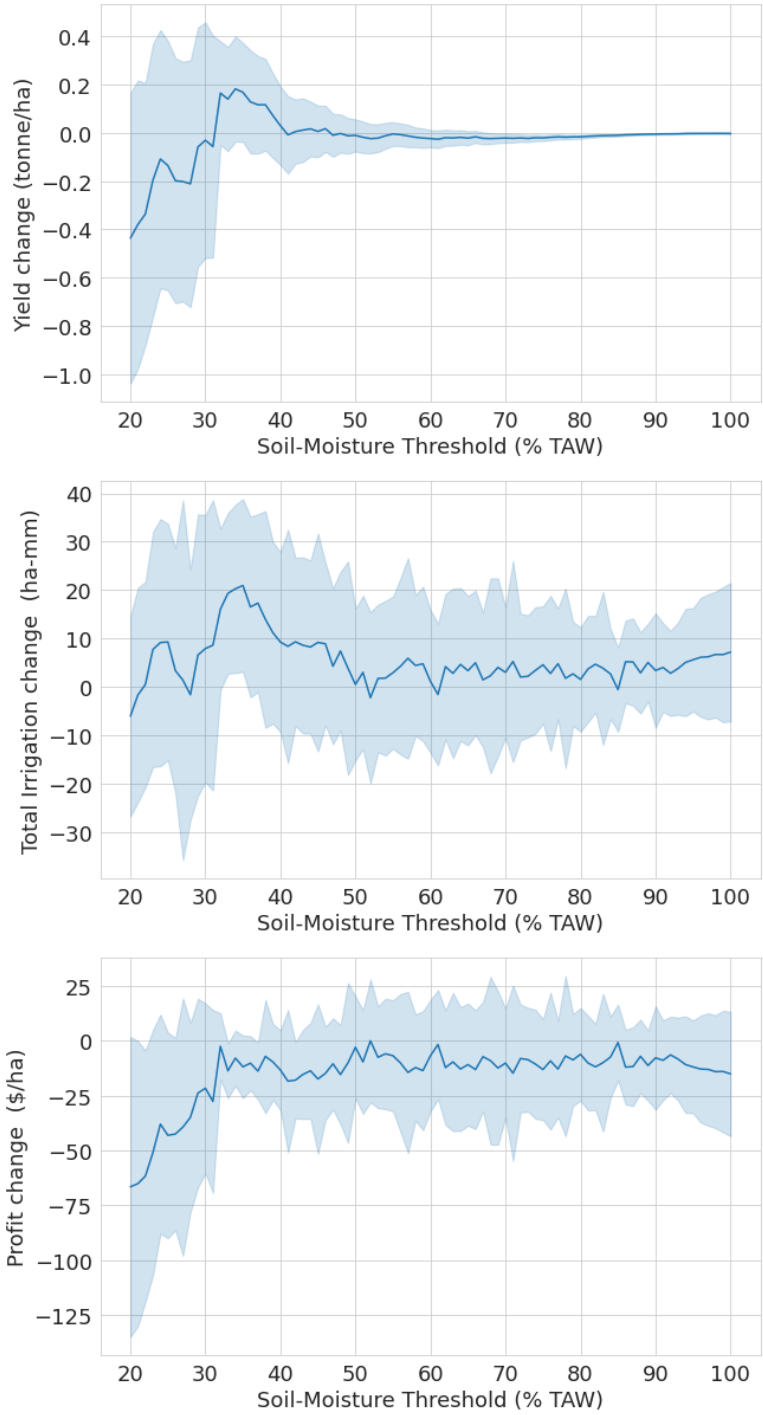

Figure 3.S7. Change in yields, water use and profits when a 30% standard error is added to water-flux and soil-texture measurements for a range of soil-moisture threshold irrigation strategies. The shaded area represents the standard deviation over 1000 30-year simulations with the solid line representing the mean over these simulations.

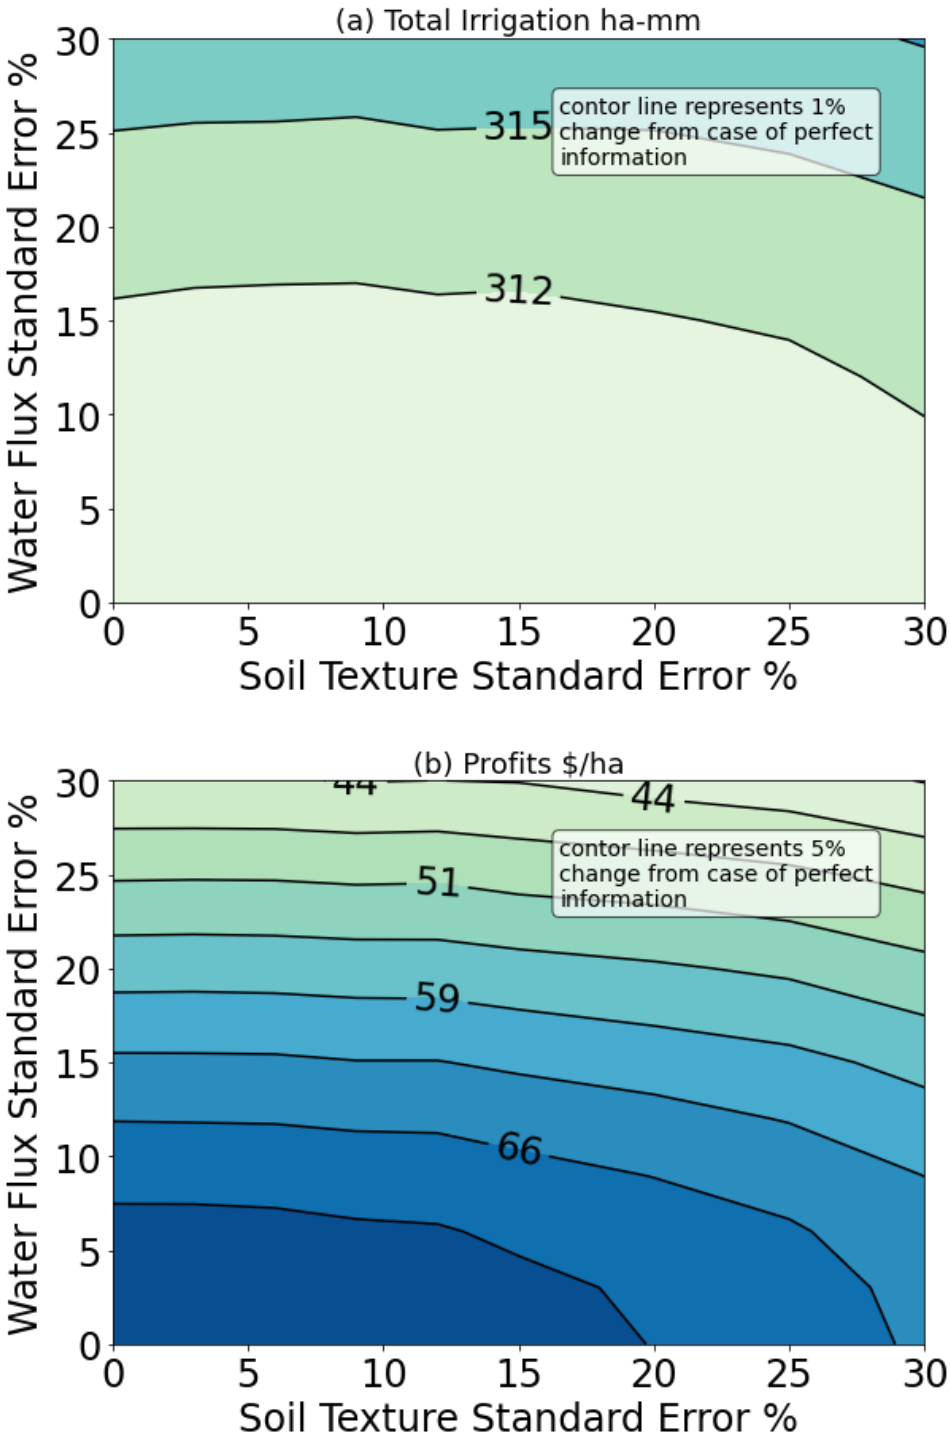

Figure 3.S8. The effect of different magnitudes of water flux and soil texture error on average total irrigation (a) and profits (b) when the crop response to water stress is turned to maximum. Each contour line represents a 1% and 5% change respectively from the case of zero added errors.

4 ASSESSING THE VALUE OF ADAPTING IRRIGATION STRATEGIES WITHIN THE SEASON

This Chapter presents the article submitted for review in May 2022:

Kelly, T. D., Foster, T., Schultz, D. M. (2022). Assessing the value of adapting irrigation strategies within the season. *Agricultural Water Management* **(in review)**

My contribution to the article was as follows: (1) Development of the research questions and methodology. (2) All programming required to run the framework and analyse results. (3) Analysing results and producing tables and figures. (4) Writing the contents of the article including literature reviews.

Assessing the value of adapting irrigation strategies within the season

T. D. Kelly[1], T. Foster[1], David M. Schultz[2, 3]

[1] Department of Mechanical, Aerospace & Civil Engineering, University of Manchester, Manchester, UK.

[2] Department of Earth and Environmental Sciences, University of Manchester, Manchester, UK.

[3] Centre for Crisis Studies and Mitigation, University of Manchester, Manchester, UK.

Corresponding author: Thomas Kelly (Thomas.kelly-2@manchester.ac.uk)

ABSTRACT

Optimization of irrigation scheduling is a widely proposed solution to enhance agricultural water productivity and mitigate water scarcity. However, there is currently a lack of knowledge about how to most effectively optimise and adapt irrigation decisions under weather and climate uncertainty, or about how the benefits of adaptive irrigation scheduling compare to fixed heuristics commonly used by farmers. In this article, we assess the added value of in-season adaptation of irrigation strategies in comparison to a fixed irrigation strategy that maximises average profits over a range of plausible weather outcomes, but is not adjusted year-to-year. For a case study of irrigated maize production in a water scarce region in the central United States, our analysis finds that fixed irrigation heuristics on average achieve over 90% of potential profits attained with perfect seasonal foresight. In-season adaptation marginally increased agricultural profitability, with greater benefits found when re-optimization occurs more frequently or is accompanied by reliable forecasts of weather for the week ahead. However, the overall magnitude of these additional benefits was small (<5% further increase in average profits), highlighting that fixed irrigation scheduling rules can be near-optimal when making realistic assumptions about farmers' potential knowledge of future weather. Since fixed irrigation strategies are easier to design, communicate and implement than data-driven adaptive management strategies, we suggest that implementing these fixed strategies be prioritised over the development of more complex adaptive strategies.

4.1 INTRODUCTION

Increasing pressures on global freshwater resources are a concern for policy-makers worldwide (Molden & FAO, 2020). As irrigation accounts for over 70% of this global freshwater use, many regions are exploring ways to reduce agricultural water use in order to mitigate current and future water scarcity and conflicts. Given the simultaneous need to increase food production to meet needs of growing populations and their changing dietary preferences, reducing or stabilizing agricultural water demands requires a focus on improving agricultural water productivity (i.e., generating more crop output per unit of water input or consumption). Strategies to achieve this goal come in the form of explicit policy restrictions on abstractions such as quotas (e.g., Leathes et al., 2008; López-Morales & Duchin, 2011; FAO, 2020; Loch et al., 2020; Young et al., 2021), as well as efforts to incentivise greater water use efficiency and productivity by improving how irrigation is scheduled and applied during the growing season (e.g., Adeyemi et al., 2017; Berbel et al., 2019; Kukal et al., 2019; Taghvaeian et al., 2020; FAO, 2020).

An extensive body of research exists that focuses on developing optimal irrigation scheduling strategies and rules to enable farmers to maximise profits or water productivity (e.g., Schütze et al., 2012; Kloss et al., 2014; Linker et al., 2016; Linker & Kisekka, 2017; Kelly & Foster, 2021). An issue faced when seeking to develop optimal irrigation management strategies is that optimal rules or decisions must hedge across multiple potential weather scenarios as farmers have at best only partial foresight of weather conditions in the upcoming growing season. A common approach in research and practice to address this uncertainty is to optimise a single irrigation management strategy that maximises the average profit over a range of potential weather outcomes such as observed in a historical record (Linker & Kisekka, 2017; Kelly et al., 2021). Alternatively, other studies have focused on selecting a single 'average' year from a collection of years (e.g., in terms of total rainfall), and optimizing irrigation management rules for that specific average year (Kloss et al., 2014).

One potential means of increasing productivity and profitability of agricultural water use would be to identify mechanisms to adapt average irrigation strategies during the season as more information is gained about weather patterns in a specific year. In this context, several studies have used adaptive simulation-optimization frameworks to evaluate potential

benefits of using weather forecasts to adapt irrigation schedules intraseasonally, concluding that such approaches can generate substantial increases in profits for producers (Wang & Cai, 2009; Cai et al., 2011; Jamal et al., 2019). Similarly, Linker (2021) also used an adaptive simulation-optimization framework to evaluate benefits of in-season re-optimization of irrigation decision rules, showing that this approach yielded irrigation schedules that were close to those that would be achievable with perfect seasonal weather foresight. However, a common limitation of these studies is that they do not assess or quantify the benefits of adaptive irrigation strategies in comparison with fixed irrigation heuristics that are commonly used by farmers and have been widely developed by researchers. As a result, these studies provide insufficient insights about the potential benefits and costs of adaptive irrigation scheduling, including potential risks of maladaptation. Better understanding of these trade-offs is critical to support improved agricultural water management decisions, in particular given the extensive data requirements and associated costs that adaptive scheduling may entail relative to use of fixed or average irrigation rules. For example, re-optimizing irrigation schedules may require additional knowledge of the current state of the crop (e.g., rooting depth, plant height, leaf-area index), which may be costly to collect alongside the time and resources required to setup and implement optimization approaches during the season.

In this study, we address the abovementioned gaps in understanding about the benefits of adaptive intraseasonal scheduling as a means of improving agricultural water use productivity and profitability under weather and climate uncertainty. We develop a simulation-optimization approach to evaluate the economic benefits of adaptive irrigation scheduling compared with fixed irrigation decision heuristics, focusing on a case study of irrigated maize production in the central United States where there are significant pressures to improve crop water productivity in response to aquifer drawdown and streamflow depletion (Scanlon et al., 2012; McGuire, 2017). Our analysis explores how the value of adapting irrigation scheduling rules within the season varies as a function of factors that have been largely neglected in past studies, including different weather scenarios, regulatory water abstraction rules, adaptation frequency and short-term weather forecasts. Our findings provide guidance to researchers, producers, and water management stakeholders on how and when adaptive irrigation scheduling impacts agricultural water use

and farm profits, and what role such approaches can play in addressing water scarcity challenges in regions experiencing chronic or growing water scarcity pressures.

4.2 METHOD

In this section, we first describe how we link a crop-growth model, AquaCrop-OSPy, to an optimization algorithm to determine a profit-maximizing irrigation strategy over a set of climate years. We then discuss how we adapt this approach to allow irrigation strategies to be re-optimised within the season to assess the added value of adaptive irrigation scheduling. We present our methods in five subsections that describe the choice of crop simulation model (Section 4.2.1), approach for identifying profit maximizing irrigation strategies (Sections 4.2.2), re-optimization of these irrigation strategies within the season (Section 4.2.3), details of the specific case study chosen (e.g., location, crop choice and model parameters) (Section 4.2.4), and an outline of sensitivity analyses performed to assess the influence of key model assumptions and production characteristics (Section 4.2.5).

4.2.1 Crop simulation model

Crop simulation models are powerful tools that simulate crop growth using biophysical equations (Jones et al., 2003; Steduto et al., 2009; Adeyemi et al., 2017). These models can be used to help design and identify strategies to enhance the productivity of water use in crop production at a much lower time and monetary cost than field experiments (Jiang et al., 2016; Araya et al., 2016; Goosheh et al., 2018). As well as the ability to compare more combinations of climate, management and irrigation scenarios, these models can be linked to optimization algorithms to more effectively discover potential scheduling strategies (Kloss et al., 2014; Linker et al., 2016; Linker & Kisekka, 2017; Kelly et al., 2021).

A wide variety of crop simulation models exist, each designed with specific use cases in mind and thus with their own strengths and weaknesses. Most come in the form of a Graphical User Interface, allowing users to alter crop, soil and management parameters and then run simulations for a given weather time series and production environment (Ragab, 2002; Jones et al., 2003; Steduto et al., 2009). AquaCrop, developed by the Food and Agriculture Organization of the United Nations (UN-FAO), is a water-driven model that aims to quantify the response of crops to water stress (Steduto et al., 2009). It contains relatively

few parameters compared with other models, allowing it to be used in data-scarce environments. Due to its popularity, AquaCrop has been successfully implemented in Matlab/Octave (Foster, Brozović, Butler, et al., 2017), R (Rodriguez & Ober, 2019), and Python (Kelly & Foster, 2021). Given the flexibility of both the model implementation (allowing us to easily integrate the model with optimization libraries) and of the Python language itself, the Python version of the model – AquaCrop-OSPy – was selected for use in this analysis.

4.2.2 Identifying optimal non-adaptive strategies

The goal of optimal irrigation management is to manage the water content in the soil to facilitate maximum crop growth with minimal water application. If the amount of soil-water available for uptake by the roots drops low enough, the crop will start to experience water stress (Ihuoma & Madramootoo, 2017). Farmers who do not use sensors or models to monitor soil-water content or crop stress precisely will still have internal heuristics about whether this 'threshold' has been reached, and therefore irrigation is required (USDA-NASS, 2018). AquaCrop-OSPy mimics this process by defining calibrated crop-water stress parameters which specify the exact level of soil-moisture that will engage the models water stress routine – reducing canopy expansion, transpiration and yields. AquaCrop-OSPy allows users to directly specify the soil-water content thresholds at which irrigation will be triggered. If the soil-moisture content drops below this threshold on a given day, irrigation will be applied to fill the soil-water content back to Field Capacity (subject to a maximum daily irrigation limit). Alternative scheduling methods to this are possible, such as specifying a regular time interval (e.g., 7 days) and depth (e.g., 10mm), irrigating to meet calculated daily evapotranspiration losses, or irrigating only once crop-water stress has been triggered. However, scheduling irrigation using direct measurements of the soil-water content – and comparing these measurements to a set of irrigation thresholds – most closely matches farmers' current practises, as well as past research (Linker et al., 2016; Foster & Brozović, 2018; Linker et al., 2018).

Given that crop-water requirements vary with growth stage (Çakir, 2004), separate soil-moisture thresholds are often defined for each stage. AquaCrop-OSPy defines four such soil-moisture thresholds corresponding to emergence, early-season canopy development, mid-season crop growth, and late-season canopy senescence. The optimal set of thresholds will

be those that maximise a particular value (e.g., profits or water productivity) for a given production environment (e.g., combination of crop type, soil type, climate conditions). This approach is similar to previous research that found profit-maximizing soil-moisture thresholds in AquaCrop either by optimization or direct search (Linker et al., 2016; Young et al., 2021; Kelly et al., 2021). For our analysis, the seasonal profit $P(s)$ is calculated at the end of each season via

$$P(s) = M * Y(s) - C * I(s) - F, \tag{4.1}$$

where *M* is a crop market price [$ per tonne], $Y(s)$ is crop yield [tonne per ha], *C* is a constant irrigation cost [$ per ha-mm], *I*(*s*) is total seasonal irrigation applied [ha-mm], and F is fixed production costs [$ per ha].

The first stage of the analysis is to define the irrigation scheduling strategy s_y (set of four soil-moisture thresholds) that maximises seasonal profit (Equation 4.1) for each weather year considered for a given case study. This process produces an optimal set of soil-moisture thresholds for each year, which we will refer to as the *potential strategy.* The resulting profits and water use from the *potential strategy* provides an estimate of what the farmer could achieve if they knew the weather perfectly in advance in all years, and adjusted their irrigation strategy accordingly based on this information. These *potential profits* therefore represent an estimated upper limit for the benefits of adapting irrigation scheduling during the season, comparable to prior studies that have optimised irrigation scheduling under perfect weather foresight (e.g., Schütze et al., 2012; Linker et al., 2016).

In reality, farmers do not know the seasons' weather with perfect foresight and irrigation scheduling decisions therefore have to be made under climate uncertainty. One way to deal with this uncertainty is to define a single strategy (i.e., a set of thresholds) that maximises the average profit across all possible years (e.g., as observed in the historic weather record). In the second part of the analysis, we identify this single average-profit maximizing strategy for our climate years by maximizing the expected (i.e., average) value of *P* over a set of N weather years. The resulting optimal soil-moisture target strategy will be referred to as the *fixed strategy.* Applying this same *fixed strategy* across all years provides a baseline for the profits that can be achieved without adapting irrigation strategies during the season. The

difference between the *potential profits* and *fixed profits* thus gives an estimated upper bound on the value of adaptive irrigation strategies.

For both the single *fixed strategy* and the multi-year *potential strategy*, the *differential evolution* global optimiser within the *scipy.optimize* python package (Virtanen et al., 2020) will be used to find the optimal set of soil-moisture thresholds. Although any algorithm can be used in principle, *differential evolution* was chosen due to its superior performance during initial tests against other optimisers within the *scipy.optimize* package.

4.2.3 Re-optimizing irrigation strategies within the season

The goal of adaptive intraseasonal irrigation scheduling is to close the gap between fixed (same strategy used for all years) and potential (perfect adjustment to each year's weather) management strategies. To assess the value of adapting irrigation strategies within season, we define an adaptive simulation-optimization framework (Figure 4.1) similar to that used in previous research (Cai et al., 2011; Jamal et al., 2019; Linker, 2021). In this framework, the simulation is paused at various points within the season, and the irrigation strategy is re-optimised from this starting point. Weather up to the current date is known with certainty as it has already been observed by the farmer, and expectations of future weather for the remainder of the season are specified based on the full ensemble of historic weather years.

We assume that irrigation decisions are re-optimised at the start of each growth stage to match the temporal disaggregation of decision heuristics commonly used by farmers and agronomists. Therefore, we first simulate crop development from planting to the end of the first crop growth stage implementing the *fixed* irrigation strategy for each weather year. At the start of the next growth stage, a new profit maximizing irrigation strategy is then found for the remainder of the season using the optimization procedure described in Section 4.2.2. The resulting strategy s is implemented over the second growth stage, after which the optimization is performed again. This process is repeated until the end of the season is reached, at which point the final profit is calculated using Equation 4.1.

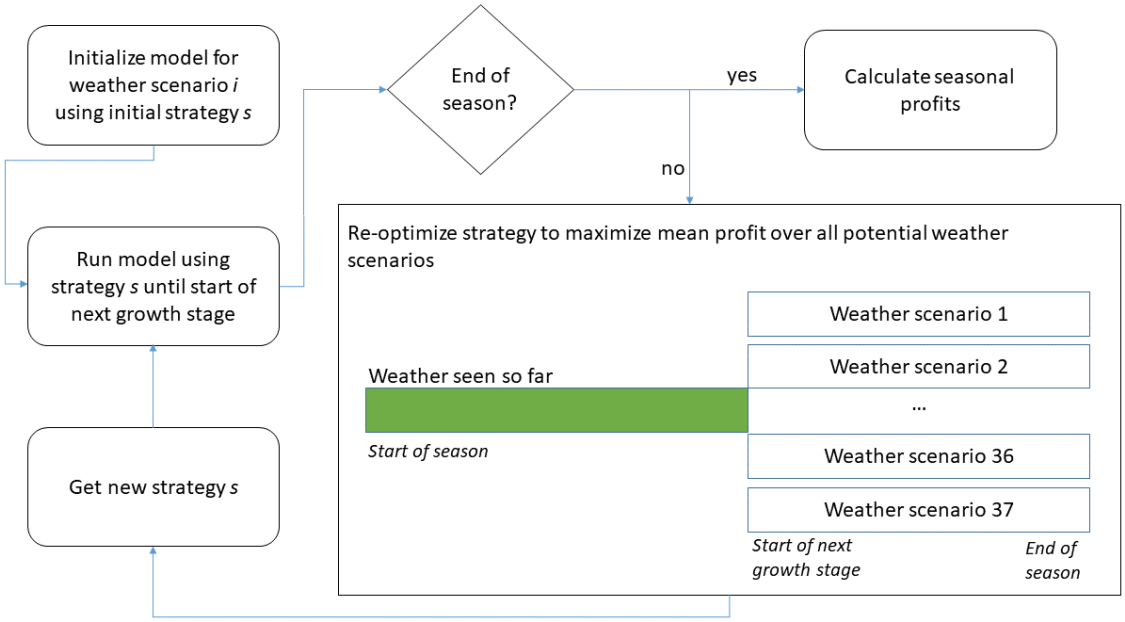

Figure 4.1. Outline of the adaptation and re-optimization framework. For each climate year, the first growth stage is simulated using the *fixed* strategy. The model's current state is then saved, and a new irrigation strategy is defined using an optimization algorithm. This optimiser finds the soil-moisture target strategy that results in the highest average profits over all the climate years, assuming the model starts at its current starting point. After finding the strategy that maximises average profit, this new strategy will be implemented for the next growth stage and the process is repeated again.

4.2.4 Case study

We apply the framework presented above to the case study of centre-pivot irrigated maize production in Nebraska, United States. Nebraska has the largest number of irrigated acres in the United States, with maize being the most dominant crop cultivated in the state (USDA-NASS, 2018). Irrigation water in the region primarily comes from the High Plains Aquifer, which has experienced drawdown over recent decades (McGuire, 2017). Associated environmental externalities, such as streamflow depletion (Szilagyi, 2000) and damage to freshwater ecosystems (Palazzo & Brozović, 2014; Perkin et al., 2019), along with concerns about long-term sustainability of groundwater-dependent rural economies (Foster, Brozović, & Butler, 2017; Deines et al., 2020; Butler et al., 2020) mean improving agricultural water productivity is a key priority for policymakers and water managers. These local characteristics make this case study an ideal choice to examine the value of adaptive irrigation decisions.

As described in Section 4.2.2, we first use the *differential evolution* optimization algorithm to find the set of soil-moisture thresholds that maximise average profit over a set of climate years. For this analysis, we use 37 years (1982–2018) of historical weather data recorded at a monitoring station in Champion, southwest Nebraska (HPRCC, 2016). The resulting optimal strategy (i.e., a set of thresholds) will be referred to as the *fixed strategy*. For each of the 37 years, the optimal irrigation strategy for each individual year (referred to as *potential strategy*) will also be calculated to provide an upper limit on what profits could be achieved in a given year with perfect seasonal foresight of weather conditions. By taking the difference between the profits achieved by the *potential* and *fixed strategy*, we can evaluate how well the *fixed strategy* performs in each year as well as the possible increases that can be made by re-optimizing within the season.

To ensure the *fixed* and *potential strategies* were as close to the global optimum as possible, 50 repetitions of each optimization were performed, and we selected the strategy achieving maximum profits across all repeat optimizations (Table 4.S1). The number of repetitions was selected to balance computational burden with optimization convergence (Figure 4.S1). During initial tests, we found that use of irrigation was never optimal during the final growth stage (i.e., crop senescence) for our case study crop and climate. To simplify the solution space for the optimization problem, this final soil-moisture threshold was therefore set to

zero for *fixed*, *potential* and adaptive irrigation strategies. As a result, only two re-optimizations were performed in the season (at the start of growth-stage two and three) when re-optimizing on a growth stage basis. The adaptation framework (Figure 4.1) for each climate year was repeated 25 times to ensure that a single sub-optimal optimization did not affect results. However, we still included the results from all of these adaption repetitions as it captures the reality that a producer will never know ahead of time what the true optimum strategy is.

AquaCrop-OSPy crop-growth parameters for maize are summarised in Foster et al. (2015), and the soil was chosen to be a clay loam, a dominant soil type in the region (CropWatch, 2018) with soil hydraulic parameters taken from the AquaCrop-OSPy default parameters. When calculating profits with Equation 4.1, a constant crop price of $180 per tonne was used based on a 10-year average of US maize grain prices (USDA, 2019). Irrigation costs were set at $1 per ha-mm ($10.28 per acre-inch) based on estimates from the 2019 Nebraska Crop Budget Report (CropWatch, 2019a). This report also provides estimates for non-irrigation production costs (e.g., labour, materials, taxes), which allows us to set the fixed production cost in Equation 4.1 to $1728 ha^{-1}. The maximum daily irrigation depth is set to be 25 ha-mm/day, which is typical for this region and production system (Young et al., 2021). Further details on model parameters can be found by viewing the full source code available at (https://github.com/thomasdkelly/adaptive-irrigation).

4.2.5 Sensitivity analysis

The value of adapting irrigation decisions within the growing season may be influenced by several factors. First, in regions experiencing water scarcity for regulatory restrictions, the amount of water farmers are able to extract for irrigation purposes is increasingly being restricted as part of water conservation and sustainability initiatives and policies. These restrictions pose issues for producers globally including our case study region. For example, the Upper Republican Natural Resource District (where our case study location resides) has imposed a five-year (2018–2022) groundwater allocation of 65 acre-inches per hectare (URNRD, 2019). These sorts of restrictions could reduce the value of adaptive irrigation strategies as producers will be less able to flexibly react to the unfolding season. We therefore repeated the main analysis (growth-stage adaptation) including a cap on the amount of irrigation that can be applied in a given season. This cap was applied in all years,

preventing any further irrigation events if the total application within the current season reached the cap. This additional analysis was performed over a range of seasonal caps from 75–550 mm ha^{-1}. For each cap, new optimal *fixed strategies* were determined, and then these were used to assess the value of adapting the strategies within the season (Table 4.S2).

Another factor that may influence the value of adaptive irrigation scheduling is the frequency with which farmers are able to re-optimise decisions. As stated in Section 4.2.4, there are only three opportunities – two when considering no irrigation is allowed in the final stage for our case study system – to re-optimise the strategy during the season within our growth-stage adaptation framework. We hypothesise that being able to adapt at more frequent intervals would lead to further increases in profits compared to the few opportunities allowed by our framework. We therefore conducted an additional analysis where our simulation-optimization framework was adjusted to allow re-optimizations to occur every seven days. Although other frequencies could also be evaluated, seven days was chosen to balance the complexity of the strategy with additional computational demand, as well as matching previous work that used similar simulation-optimization frameworks (Cai et al., 2011; Jamal et al., 2019). The first re-optimization will therefore occur on day 7 of each season and re-optimizations will stop after day 120, resulting in 17 different sets of thresholds being implemented throughout the season. Note that the choice of day 120 of the season as the final date for re-optimization was because this day falls after the start of the final growth stage (when irrigation is always zero), but this date could be varied to any other value in alternative model case studies or applications.

Finally, as noted earlier, the ability to effectively update irrigation decisions may also be influenced by the quality of information available to a farmer about weather conditions in the days or weeks ahead. In our main analysis, re-optimization is performed over the same set of weather years used in the *fixed* strategy. Producers must therefore hedge the irrigation strategy against all 37 potential weather scenarios for the remainder of the season. In reality, some of these years will be more probable than others given weather so far and available forecasts. The seven-day adaptation framework described above was therefore further adjusted to include a seven-day perfect forecast. Beyond those seven days, the producer must again hedge across all 37 weather years. The seven-day lead time

for the perfect forecast was chosen to match previous work (Jamal et al., 2019). We hypothesise that including the seven-day perfect forecast will increase profits compared to the standard seven-day adaptation due to the greater ability to tune irrigation decisions to the weather ahead for a given season.

The experiments were performed on The University of Manchester's High Performance Computing (HPC) cluster with each job being computed on a 2×8-core Intel Xeon E5-2650 v2 @ 2.60GHz + 64GB RAM compute node. Up to 20 compute nodes could be used simultaneously, each running a separate years simulation. To run growth stage optimisation on all 37 years could be completed in approximately one day, and hence up to a week to perform the analysis for all irrigation caps. Running seven-day re-optimization experiments would take approximately three days to complete.

4.3 RESULTS

4.3.1 Differences between *fixed* and *potential* strategies

The first part of this analysis used the *differential evolution* optimization algorithm to find a set of soil-moisture thresholds that maximised average profits over 37 climate years, with the resulting set of thresholds being referred to as the *fixed strategy*. The *potential strategies* (i.e., profit-maximizing strategies for each individual year) were also calculated to give an estimated upper limit for what profits could be achieved with perfect seasonal information (Table 4.S1). A comparison between the average profits achieved by the *fixed* ($426 ha^{-1}) and *potential* ($459 ha^{-1}) strategies finds that the *fixed strategy* achieved 92.7% of the average *potential profit* (Figure 4.2a). In other words, having perfect foresight of weather conditions in each year results in a 7.9% increase in farm profits.

Looking more closely at each individual climate year, there is only one year (2012) in which the *fixed strategy* does not achieve at least 80% of the *potential profits* (Figure 4.2a). This year was one of extreme drought (less than 43 mm of rainfall during the growing season) and so appears to be sacrificed by the optimiser in favour of maximizing average profits across all remaining years. Focusing instead on the other 36 non-drought years, regression analysis indicates that 40% of the variation in the performance of the *fixed strategy* can be

explained by the total seasonal rainfall (Figure 4.2b). The *fixed strategy* performs better (in terms of percentage of *potential profits*) in the wetter years than the drier years. Using the *fixed strategy* as a starting point, the framework described in Section 4.2.3 aims to assess whether adapting the strategy within the season will increase these *fixed profits* towards the *potential profits*.

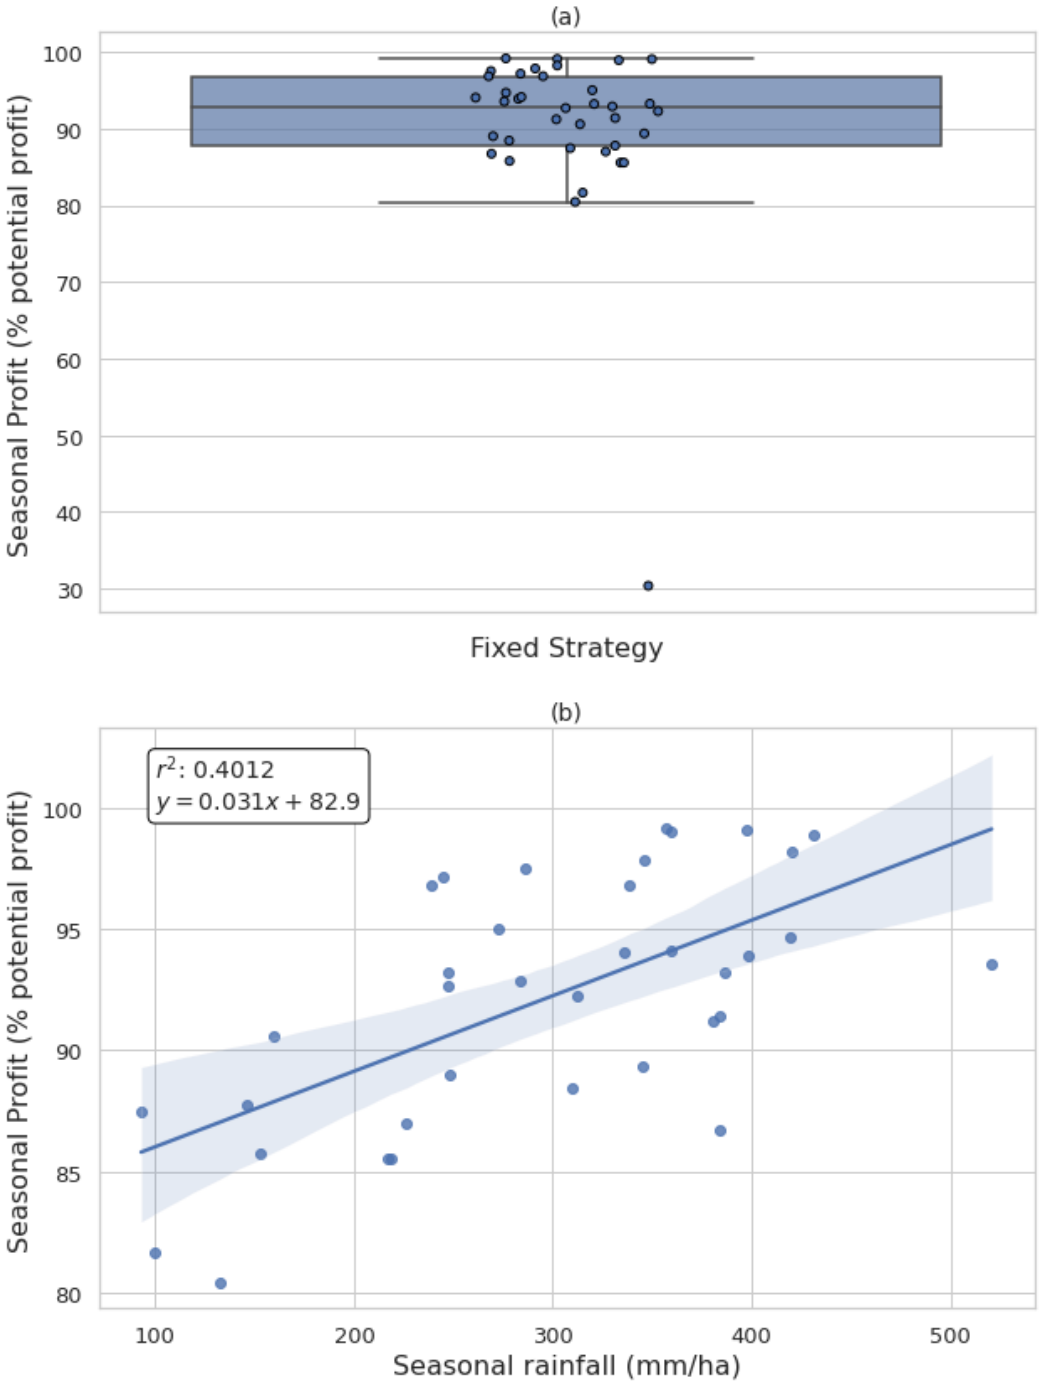

Figure 4.2. (a) Seasonal profits achieved but the *fixed strategy* for each climate year expressed as a percentage of the *potential profits* for that year. (b) Seasonal profits (% potential profits) for each climate year (excluding the year 2012) as a function of total rainfall during that growing season. Comparable results (not shown) were also obtained when using the seasonal water balance (precipitation minus crop evapotranspiration) as a predictor.

4.3.2 The value of within-season adaptation

To assess the value of re-optimizing irrigation strategies within the season, the *fixed* irrigation strategy (set of soil-moisture thresholds) was re-optimised at the start of each growth stage following the re-optimization approach described in Section 4.2.3. After running this framework for all 37 climate years, the profits from this adaptive strategy can be compared to the *fixed profits* to determine the added value of this adaptive approach to scheduling.

On average, our analysis found a \$4.87 ha^{-1} increase in profits with adaptation, which translates to a 1% increase in average profit compared to the *fixed strategy* that does not incorporate any adaptation of irrigation management rules during the growing season. However, we also observe that growth-stage adaptation also increased the minimum profits by 106% and decreased the standard deviation in profits by 2.95%. The large increase in minimum profits can be attributed to the drought year 2012, where in-season adaptation increases the soil-moisture thresholds, causing additional irrigation events and profit increases. By inspecting the irrigation schedule time series for this year, we see the impact of this change in soil-moisture thresholds on the soil-water content and the timing of irrigation events (Figure 4.3). Changes to the minimum profits and profit variance may have important implications for some farmers, in particular where there is a need to minimise risks associated with extremely low yields in individual years as opposed to simply looking to maximise average profits over multiple years. These implications will be further discussed later in Section 4.4.1.

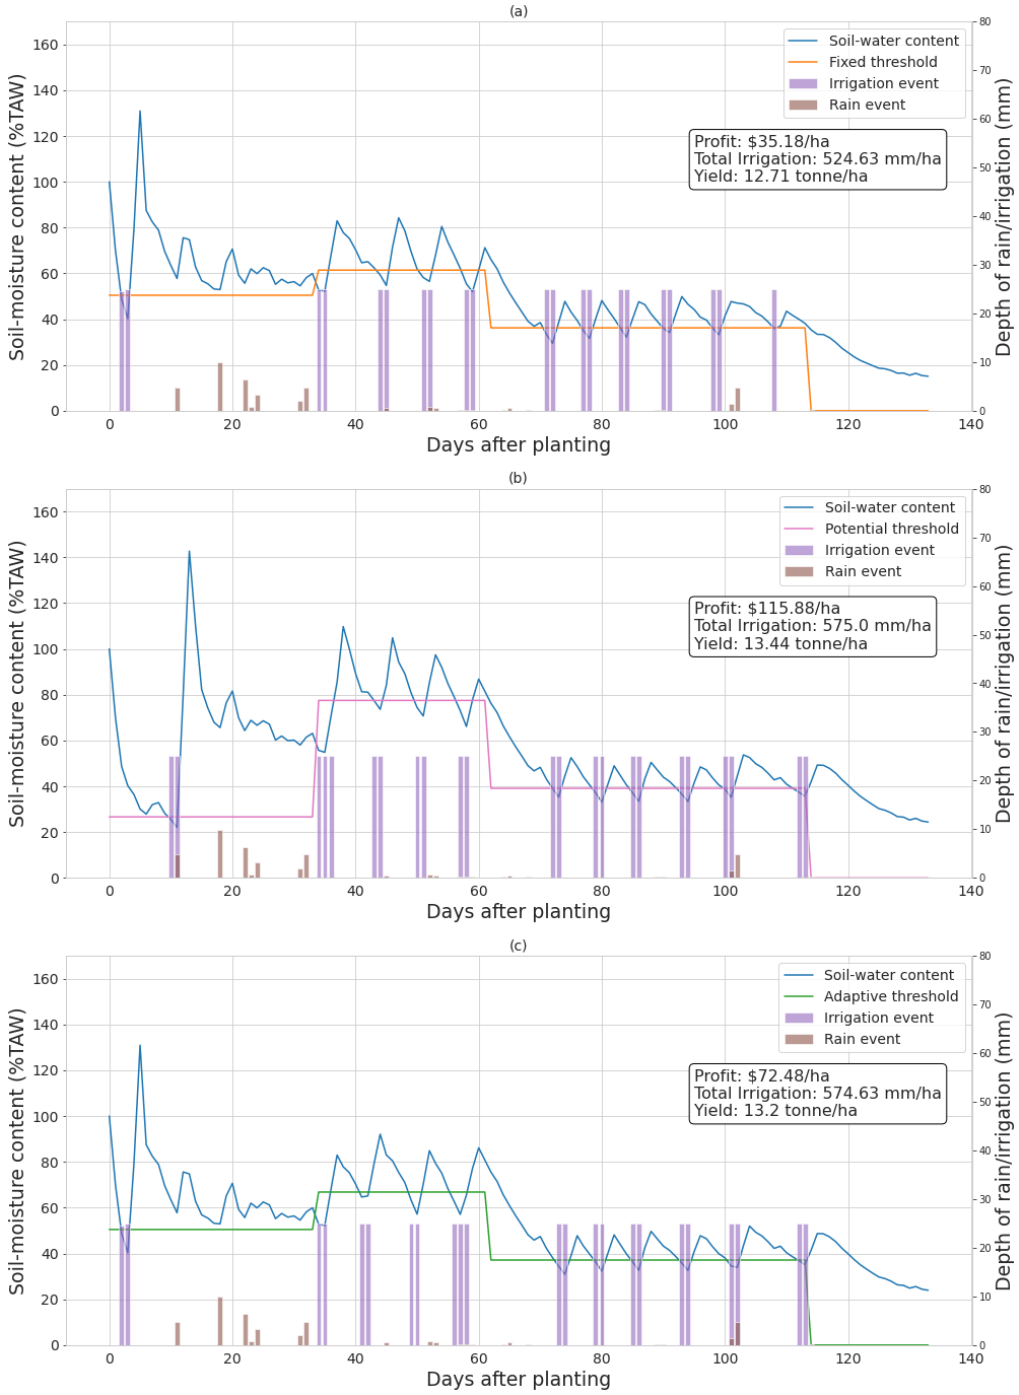

Figure 4.3. Soil-moisture content and irrigation thresholds across the 2012 growing season for (a) *fixed strategy*, (b) *potential strategy*, (c) growth-stage adaptation strategy. When the soil-moisture content drops below the irrigation threshold, irrigation is triggered. The depth of irrigation and rainfall events are shown in purple and brown bars respectively. This season represents one of the 25 repetitions.

Several factors are important drivers of the value added from in-season adaptation of irrigation scheduling rules. First, we find that the quality of the starting strategy – defined as how close the *fixed strategy* mean profit is to the *potential strategy* in a given year – explains 39% of the variation in adaptation value (Figure 4.3a). This relationship shows that adaptation value is lower in years where the *fixed strategy* is already close to the potential optimal for that year. In addition, by assessing the relationship between seasonal rainfall and adaptation value, we find that adaptation value is inversely proportional to the total seasonal rainfall. However, our analysis indicates that seasonal rainfall only explains 14% of the variation in adaptation value (Figure 4.3b). As well as total rainfall, how the rainfall is distributed over the season is also shown to impact optimal irrigation strategies. Specifically, we find that the coefficient of variation (CV) in daily rainfall (i.e., standard deviation in daily rainfall / mean daily rainfall) explains 19% of the variation in adaptation value (Figure 4.3c). Both of these relationships indicate that when rainfall is scarce or more variable, the value of within-season adaptation is greater.

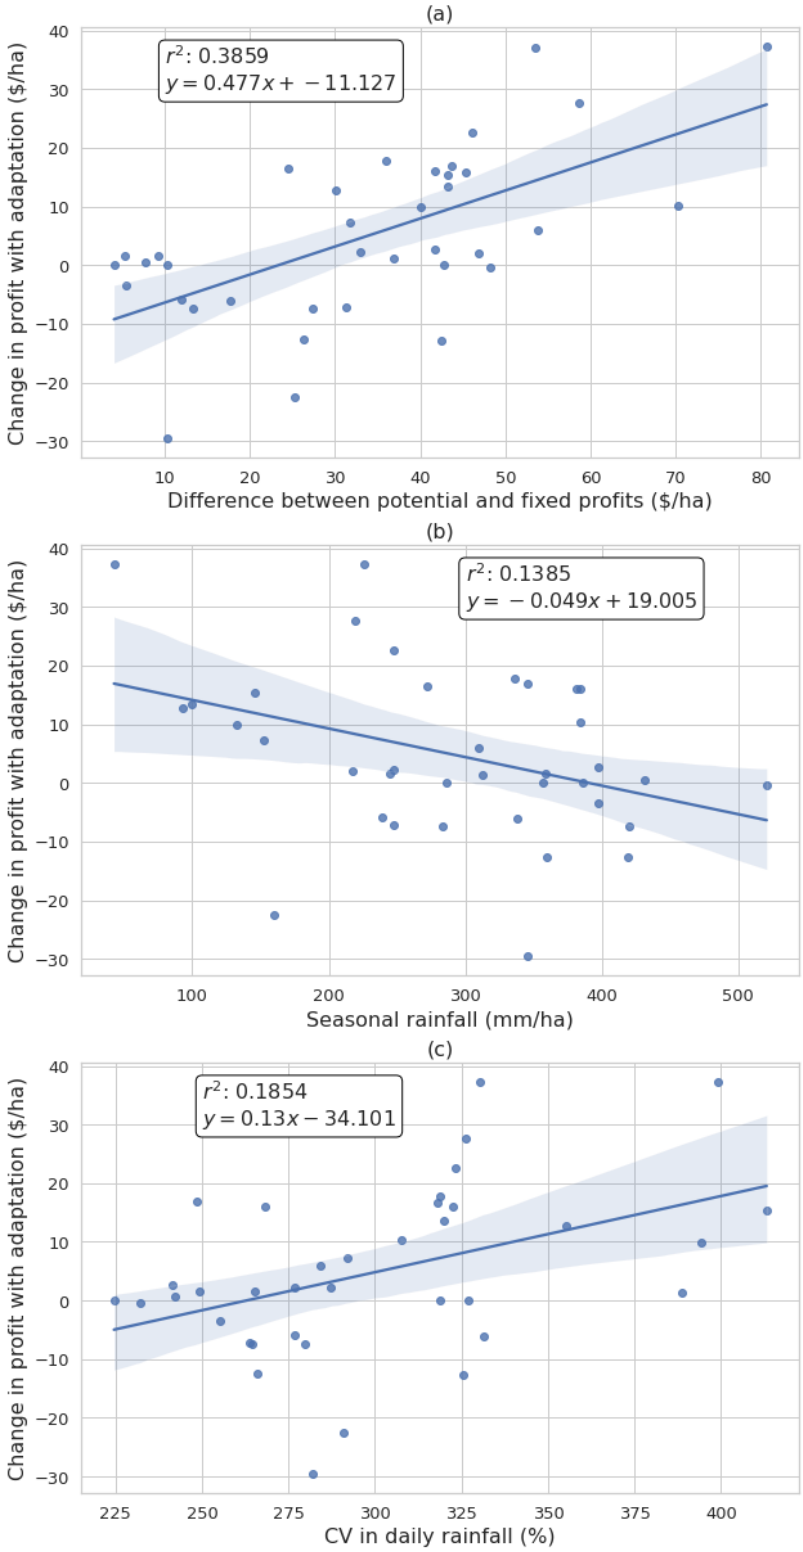

Figure 4.4. The relationship between: (a) adaptation value and quality of starting strategy, (b) adaptation value and total seasonal rainfall, (c) adaptation value and CV in daily rainfall. Each point represents one of the 37 climate years, and adaptation value is averaged over 25 repetitions of the analysis framework. Shaded areas represent the 95% confidence interval for the linear regression.

Inspecting irrigation schedule time series for specific individual years also provides additional insights about why in-season adaption is beneficial in some years and not in others. For example, Figure 4.5 compares the irrigation events triggered by the *fixed* (Figure 4.5a), *potential* (Figure 4.5b) and growth-stage adaptation (Figure 4.5c) strategies during the 1999 growing-season simulation, where adaptation decreased seasonal profits by \$30 ha^{-1} (the largest decrease for any year). Overall, during both the second and third growth stage, the adaptive strategy made minor adjustments to the soil-moisture threshold towards the *potential* threshold. One would expect that such an adjustment (towards the perfect information strategy) would increase profits – as it does for the year 2012 – however this does not occur in the year 1999. This slightly counterintuitive result can likely be attributed to having to optimise decisions under future weather uncertainty, as one can never know what the exact optimal set of thresholds are until the entire season has passed. Changes to the soil moisture thresholds are unable to perfectly replicate irrigation schedules under the potential strategy with perfect foresight, with two extra irrigation events being triggered on days 89 and 90 for the adaptive strategy. This reduces overall profits for the season, demonstrating that, without perfect knowledge of the full seasons' weather beforehand, there will always be the possibility for maladaptive behaviour (e.g., the adjustments made in 1999) during some years.

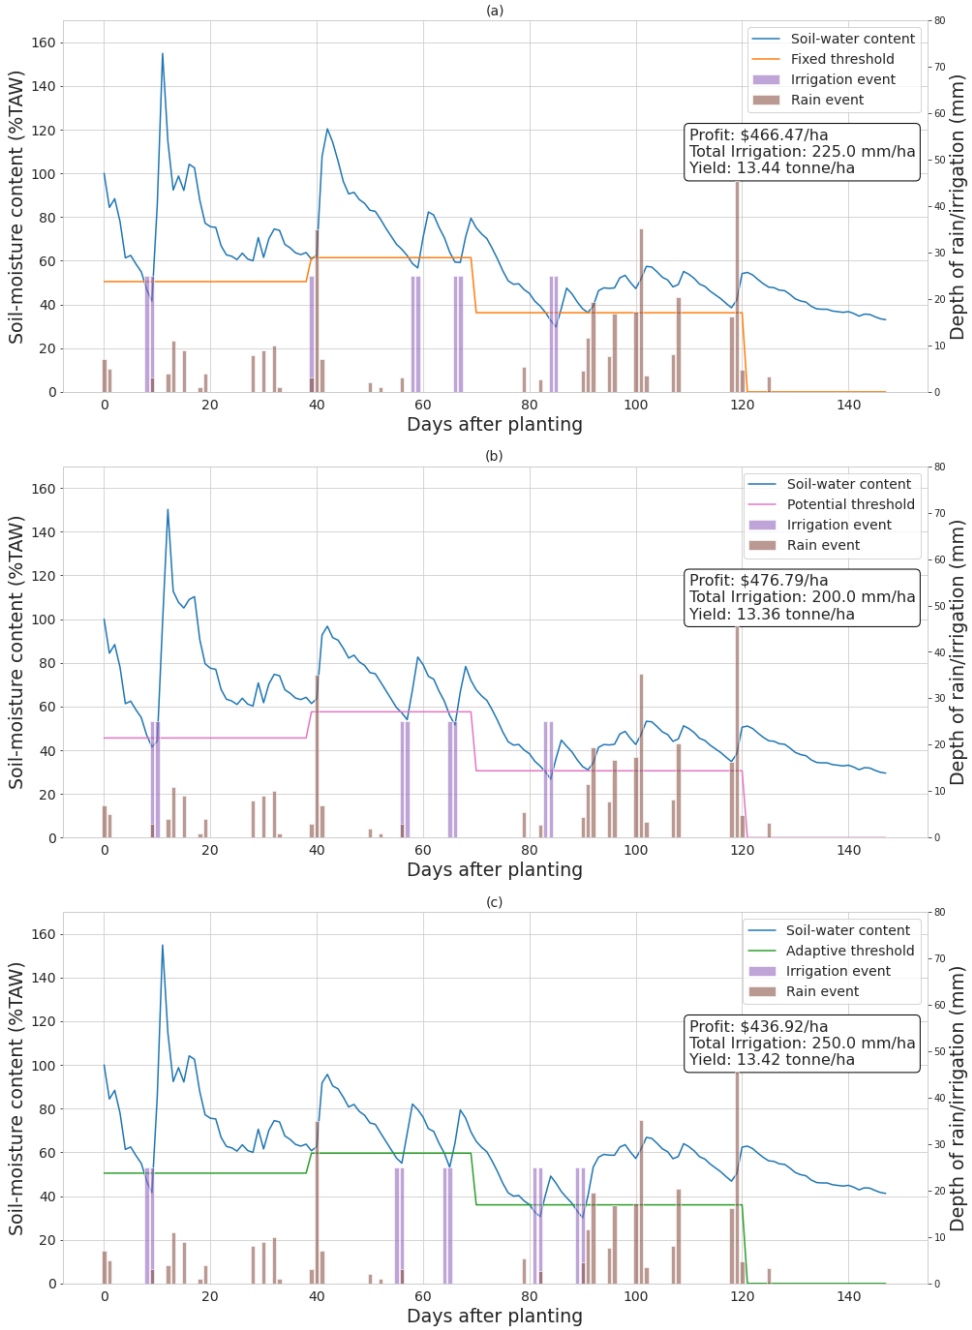

Figure 4.5. Soil-moisture content and irrigation thresholds across the 1999 growing season for (a) *fixed* strategy, (b) *potential* strategy, (c) growth-stage adaptation strategy. When the soil-moisture content drops below the irrigation threshold, irrigation is triggered. The depth of irrigation and rainfall events are shown in purple and brown bars respectively.

4.3.3 The impact of adaptation frequency and forecast information on adaptation value

As described in Section 4.2.5, two further analyses were conducted to test whether adaptation frequency and quality of forecast information impacted the value of re-optimization. Allowing re-optimization of irrigation strategies every seven days lead to a \$2.25 ha^{-1} increase in profits compared to growth-stage re-optimization, achieving 93.69%

of the *potential profits* over the 37 climate years (Figure 4.6a). By including a perfect seven-day weather forecast within this seven-day re-optimization framework, average profits were increased by a further \$13.4 ha^{-1} compared to the standard seven-day re-optimization, achieving 97.28% of potential profits.

These analyses confirmed our hypothesis that re-optimizing more frequently, in particular when combined with better information about weather conditions in the days ahead, increases the benefits that can be derived from in-season adaptation. As well as changes to average profits, we also see changes in profit variance with adaptation. Switching from the growth-stage adaptation to seven-day adaptation increases the minimum profits by 15.8% and reduces the standard deviation in profits by 2.56% compared to growth-stage adaptation (Figure 4.6b). Adding a perfect seven-day weather forecast to the seven-day adaptation strategy further increases the minimum profits by 29.2% compared to the seven-day adaptation strategy without perfect weekly forecasts.

An additional notable effect of introducing perfect weekly forecasts to the seven-day adaptation strategy is that there are eight years where the seven-day adaptation with perfect weekly forecast outperforms the *potential strategy*. This result indicates the limitations of how we have chosen to define the *potential strategy*, which – despite having perfect seasonal weather information – still only contains one fixed set of thresholds for the season. In comparison, the seven-day adaptation strategy contains 17 sets of thresholds that are calculated as the season progresses. The ability to vary soil moisture targets within – rather than just between – growth stages allows for more complex irrigation management strategies to emerge, sometimes resulting in higher profits despite the farmer not having access to perfect information about weather for the entire season.

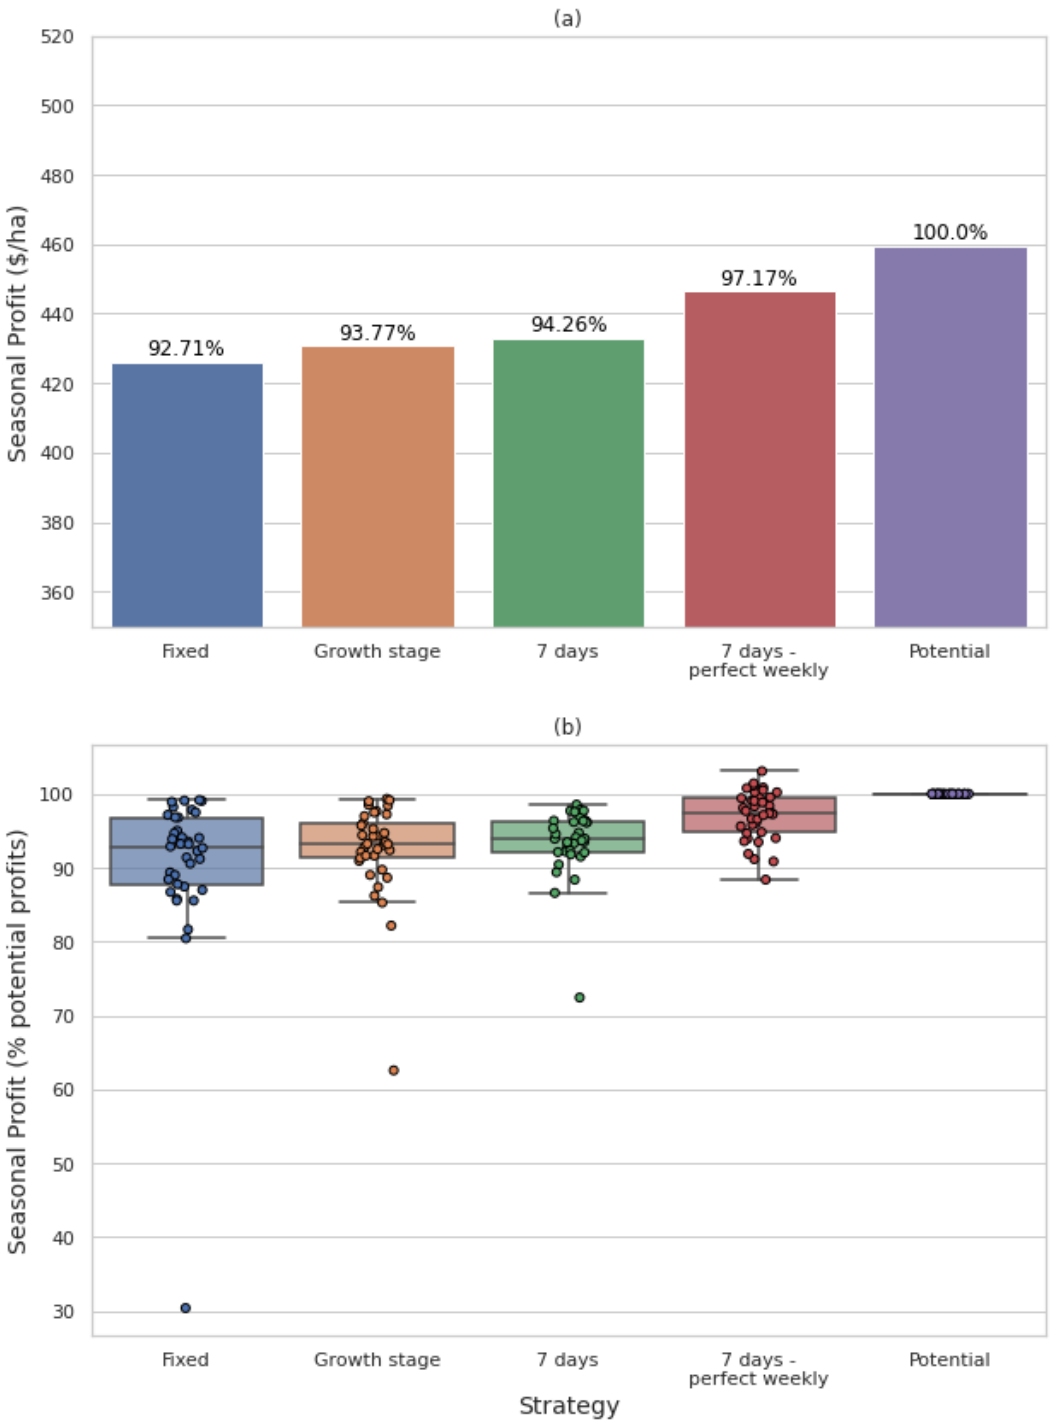

Figure 4.6. (a) Seasonal profits averaged over 37 climate years for each adaptation strategy. This quantity is also displayed as a percentage of average *potential* profits. (b) Seasonal profits for each climate year expressed as a percentage of the *potential* profits for each adaptation strategy.

To demonstrate an example of when the seven-day strategy with perfect weekly forecast outperformed the *potential strategy*, we compare the individual irrigation events triggered by the *fixed* (Figure 4.7a), *potential* (Figure 4.7b) and seven-day adaptation with perfect weekly forecast (Figure 4.7c) strategies for the climate year 1992. The flexibility enabled by the seven-day adaptation period initially allows an increase in the first growth-stage threshold (Figure 4.7c), meaning that irrigation events for the *potential* and adaptive strategies were identical for the first 45 days (Figure 4.7b, 4.7c). After this initial period, the two strategies diverge as the adaptive strategy sticks closer to the *fixed strategy* (with the flexibility to adjust as required), whereas the *potential* thresholds diverge massively during the second growth stage in particular.

Although the seven-day adaptation strategy with perfect weekly forecast occasionally outperformed the *potential strategy*, there are still a handful of years where it also performed worse than the *fixed strategy* (Figure 4.S2). As stated in the previous section, this effect is the result of having to optimise irrigation strategies under uncertainty about weather beyond the seven-day forecast, and is not the result of the optimization approach or choice of strategy. However, switching from the growth-stage adaptation to the seven-day adaptation increased does profits in these years. This result implies that more complex irrigation strategies that also incorporate more accurate weather information can help reduce the risk of profit losses from these maladaptive events, even when knowledge of weather conditions for the remainder of the season remains unknown or uncertain.

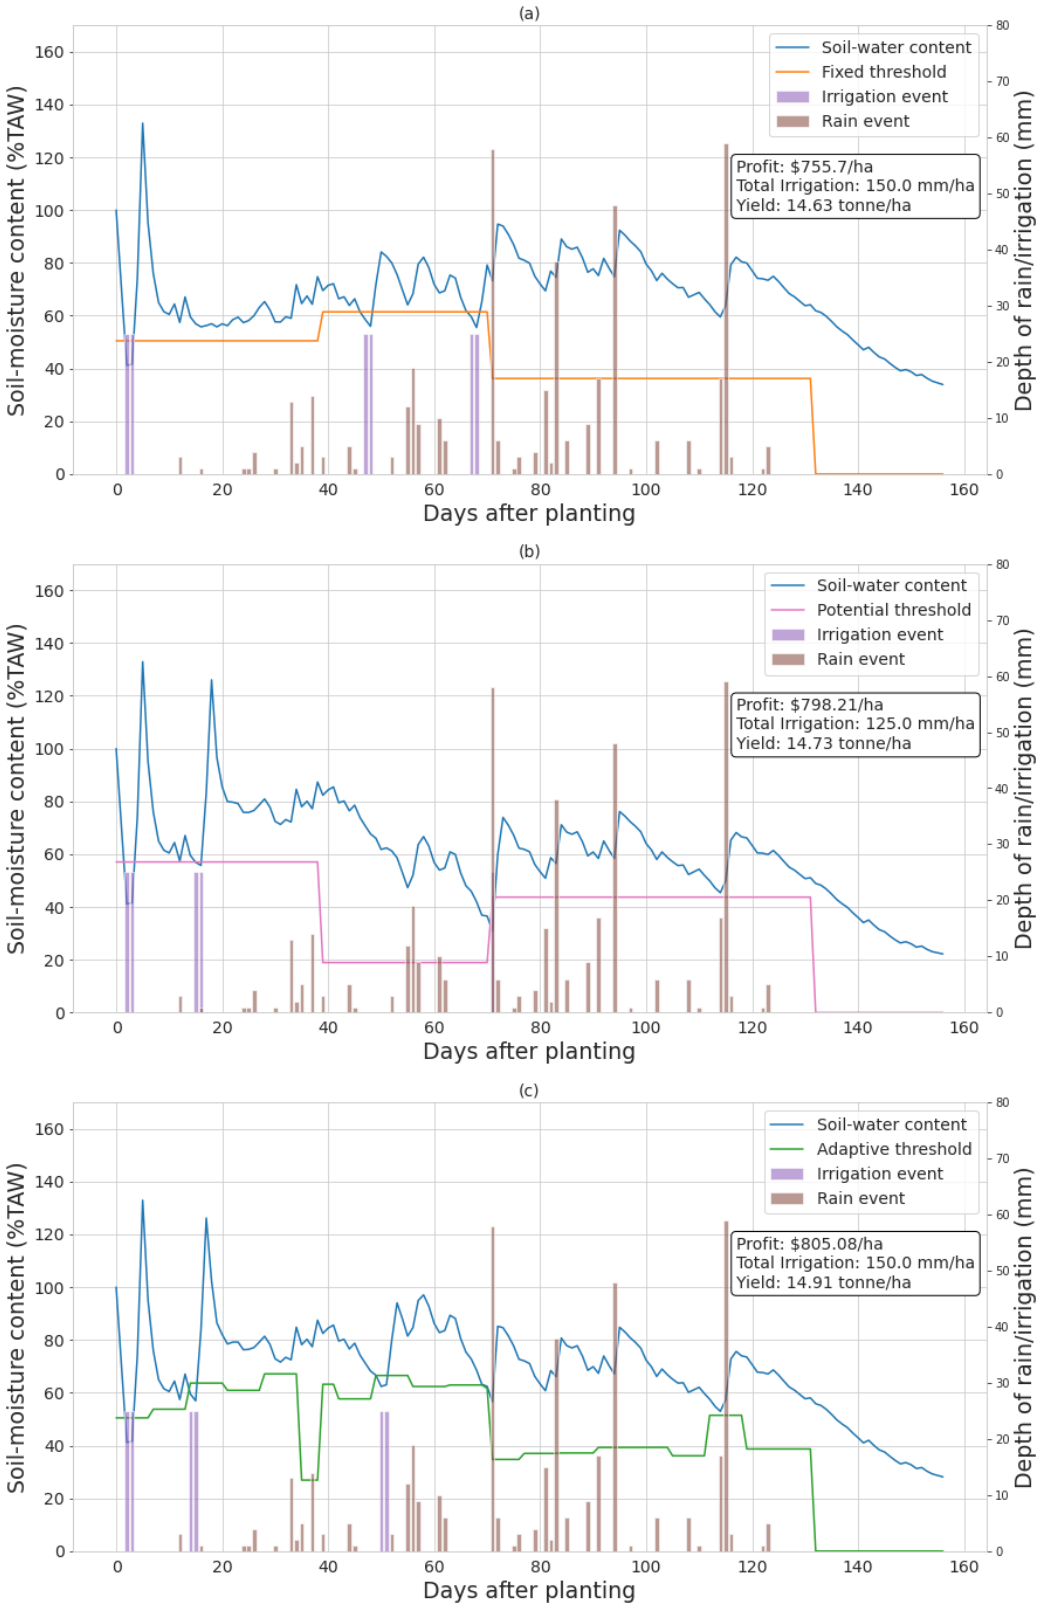

Figure 4.7. Soil-moisture content and irrigation thresholds across the 1992 growing season for (a) *fixed* strategy, (b) *potential* strategy, (c) seven-day adaptation strategy with perfect weekly forecast. When the soil-moisture content drops below the irrigation threshold, irrigation is triggered. The depth of irrigation and rainfall events are shown in purple and brown bars respectively.

4.3.4 The impact of water constraints on adaptation value

Regulatory or hydrologic restrictions on irrigation water supply are a common consequence of water scarcity and conflict in many regions, including our study area in the central United States. We therefore explored how the value of adaptation was affected by the introduction of annual caps on total seasonal irrigation that, once reached, would prevent any further irrigation being applied during the season.

By introducing a seasonal irrigation cap of 300 mm ha^{-1}, average seasonal profits achieved by the fixed strategy decrease by \$82 ha^{-1} in comparison to the \$426 ha^{-1} achieved when water use was unrestricted. As annual water quotas increase, profits increase towards the unrestricted fixed profits. For the 300 mm ha^{-1} limit, the average value of adaptation was \$1.49 ha^{-1} (0.4% increase from fixed strategy). This value increases linearly with seasonal water allocation reaching a benefit of \$4.87 ha^{-1} (1% increase from fixed strategy) achieved when water use is unrestricted (Figure 4.8a). These results imply that restrictions on yearly water use actually decrease the value of adaption. A possible intuition for this result is that having more water available during the season allows greater flexibility for the re-optimization process to adapt within the season.

For more severe restrictions on water use (i.e., quotas less than 300 mm ha^{-1}), impacts of adaptation begin to vary significantly from year to year as well on average, with large benefits in some years and large costs in others (Figure 4.8b). These results show that while some benefits may still be derived from adaptation under highly restrictive water use limits, the variation in outcomes and potential for maladaptive responses is large. In contrast to the analyses without water use restrictions presented previously (Figure 4.4), we find no relationship between adaptation value and weather conditions under severe water use restrictions. Specifically, the r^2 between adaptation value and total rainfall for the most restrictive caps (75, 100, 125 mm ha^{-1}) were all found to be less than 0.02. Similarly, the r^2 between adaptation value and CV in daily rainfall for these caps were also found to be less than 0.02. Therefore, whether or not adaptation leads to extreme benefits or costs appears to be mostly down to luck for these extreme water use limits.

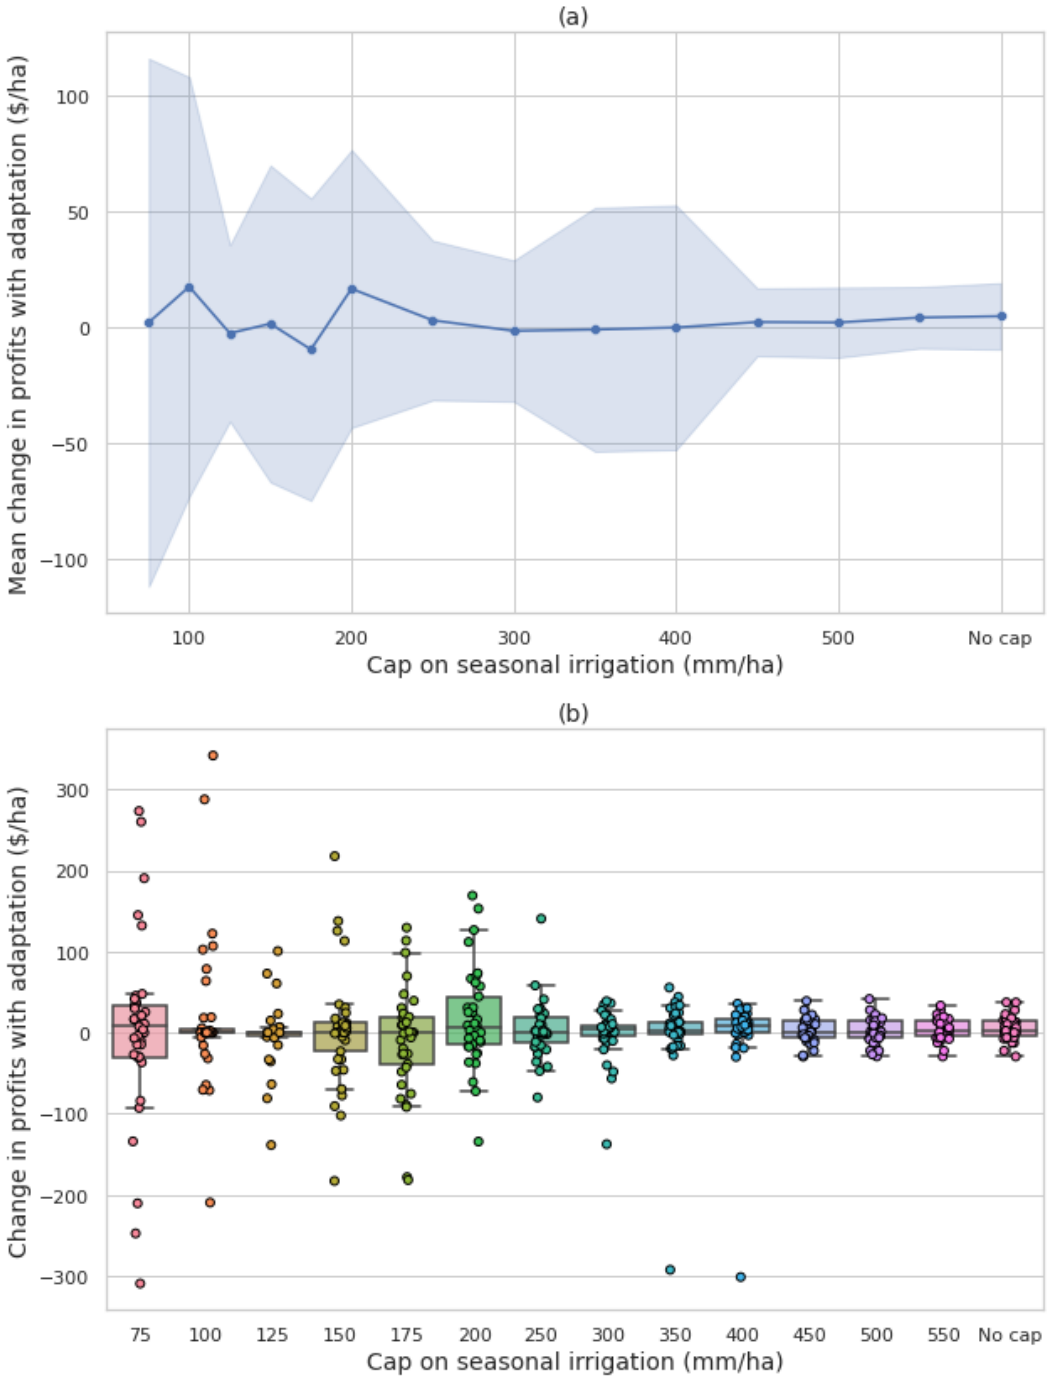

Figure 4.8. (a) Change in mean seasonal profits with adaptation (i.e., adaption value) over the 37 climate years for irrigation caps ranging from 75 mm ha^{-1} to unconstrained irrigation. Blue shaded area indicates standard deviation between years. (b) Distribution of adaption value between years for each irrigation cap. Each point represents a single year averaged over multiple simulation repetitions.

4.4 Discussion

Optimization of irrigation scheduling, including through in-season adaptation of scheduling rules based on weather forecasts and other data, have been widely proposed as a means of improving agricultural water productivity in water-scarce regions globally (e.g., Wang & Cai, 2009; Cai et al., 2011; Hejazi et al., 2014; Jamal et al., 2018, 2019; Linker, 2021). By linking a crop model to our re-optimization framework, this article aimed to assess the value of adapting irrigation strategies within the season, and how this value is affected by typical information constraints (e.g., in-season weather) and management constraints (e.g., abstraction restrictions) commonly faced by farmers but not adequately considered in past research (e.g., Wang & Cai, 2009; Cai et al., 2011; Jamal et al., 2019).

For a case study of maize production in a water stressed region of south-west Nebraska in the United States, our analysis demonstrated that a farmer adopting a *fixed* irrigation strategy (i.e., where rules are the same for each year) was able to achieve 92.7% of the average *potential profit* achievable with perfect seasonal weather foresight. In-season re-optimization of *fixed* irrigation strategies based on information gained about weather conditions as the seasonal unfolds was able to marginally increase average profits to 93.8% or 94.3% of *potential profits* depending on how frequently farmers choose, or are able to, adapt irrigation strategies within the season.

We demonstrate that value of in-season adaptive re-optimization of irrigation heuristics is enhanced when farmers have access to perfect weather forecasts for the week ahead, for which the producer is able to achieve 97.2% of potential profits. Adaptive scheduling also universally reduced risks of crop loss or failure events (e.g., during extreme drought years such as 2012 in our study area), but with greatest risk reduction benefits found when re-optimization is supported by in-season forecasts. By introducing seasonal water quotas, our analysis showed that the variability in adaptation value between years increased dramatically for extremely restrictive water quotas.

4.4.1 Implications for research, producers and industry

Our results build upon previous work that aimed to optimise irrigation strategies within crop models using adaptive re-optimization during the growing season (e.g., Wang & Cai, 2009; Cai et al., 2011; Hejazi et al., 2014; Jamal et al., 2018, 2019; Linker, 2021). These works did not include an optimised fixed strategy with which to compare the value of adaptive irrigation strategies. As a result, these studies are likely to have overstated the true potential benefits of adaptive in-season irrigation scheduling, with our results demonstrating that the majority of potential profits can be achieved with simpler optimal average irrigation heuristics. Our findings are also consistent with other evidence about the value of adaptive agricultural decision making in the literature. For example, Jones et al. (2000) showed that optimization of crop management practices for maize (e.g., planting date, crop density, nitrogen application) based on seasonal climate forecasts increased average seasonal profits by only \$16–26 ha^{-1} (3–4%) depending on the case-study region when compared with fixed management rules based on a range of historic weather years. Our findings thus highlight the importance of comparing adaptive strategies to similarly optimised fixed strategies, which may provide a more efficient means of scheduling irrigation when technological, logistical and cognitive costs associated with implementing more complex data-driven adaptive scheduling practices are taken into account.

A key finding of our analysis is that there is likely to be substantial heterogeneity in the economic benefits of adaptive irrigation scheduling across years, with some years potentially experiencing declines in profits when strategies are re-optimised within the season. This type of maladaptation was also found in Hejazi et al. (2014) and occurs as a result of underlying uncertainty about future weather, along with the complex non-linear responses of crop growth and yield development to soil-moisture patterns during the season. Unless the entire season's weather is known with a high degree of certainty in advance, it is extremely challenging for a farmer to determine which irrigation strategy will reliably lead to the largest profits. Even if one strategy is superior on average, there will be some years where an alternative strategy results in higher profits. The season-to-season variability in profits also highlights that instead of evaluating irrigation strategies on as few as 5 climate years – as is common in previous work (Wang & Cai, 2009; Cai et al., 2011; Hejazi et al., 2014; Jamal et al., 2018, 2019) – researchers must use quantities approaching

30 in order to fully capture local weather characteristics and thus accurately evaluate outcomes of adaptive decision-making or the value of data (e.g., forecasts).

Our analysis found that some of the year-to-year variability in benefits from adaptive irrigation scheduling could be explained by the quality of the starting strategy, total rainfall during the season, and the coefficient of variation of daily rainfall during the season. These relationships suggested that in years with less rainfall, or where rainfall is unevenly distributed during the season, the economic value of adaptive scheduling will be higher. This is particularly relevant given climate change is predicted to increase the frequency of weather extremes in many regions (IPCC, 2022). Our analysis also found that adaptive strategies did cause an increase in minimum profits and decrease in profit variance over the 37 years. For more risk-averse farmers who cannot simply seek to maximise average profits, this reduction may provide enough motivation to invest in adaptive strategies. However, the presence of other safety nets, for example crop insurance schemes, could decrease the value of adaptive scheduling practices as farmers will be financially protected against some of the consequences in the driest years.

The strong performance of both fixed and adaptive irrigation strategies relative to theoretical *potential profits* is in part due to specifying a strategy made up of four soil-moisture thresholds, one for each major growth stage. Kelly et al. (2021) found that choosing a simpler strategy made up of just one soil-moisture threshold for the entire growing season led to a \$120 ha^{-1} decrease in profits compared to the four soil-moisture threshold strategy as done in our analysis. This decrease is three times larger than the profit difference found between the *fixed* and *potential* strategies in our analysis. Crop models and simulation-optimization frameworks such as that used in the present article are vital for designing and optimizing these irrigation strategies, which can also be replicated for many crops and production environments. Over time, however, these optimal strategies will need to be updated to ensure the historical weather years used to optimise the strategy reflect changes in the likely range of weather scenarios that will occur in the upcoming season. For example, changes in the variability of expected weather outcomes due to climate change will likely reduce the performance of fixed strategies and enhance benefits of adaptive scheduling if these shifts in climate variability are not accounted for when designing fixed irrigation management rules.

4.4.2 Framework Limitations and Future Directions

While our analysis provides valuable insights about the potential value of adaptive in-season irrigation scheduling for improving agricultural productivity and profitability, it is important to highlight several key simplifications in our analysis.

First, an assumption made in our analysis is that the re-optimization framework considers the hypothetical farmer to have perfect knowledge of the current state of the soil water balance and crop growth (e.g., yield, canopy cover, rooting depth). However, in the real world, farmers may not have perfect information about current field conditions, or be able to predict with certainty how crop development will be affected by different potential weather scenarios. As a result, the real-world value of within-season adaptation is likely to be even lower than found in our analysis, as the re-optimization of irrigation strategies will be influenced by additional factors. These factors include the uncertainty in weather conditions for the days and weeks ahead, as well as uncertainties in the current status of crop growth and soil moisture. These additional uncertainties are likely to increase the potential for maladaptive irrigation scheduling responses to occur, in particular where capacity for irrigation water use is limited by regulatory, hydrologic, or socio-economic constraints. However, further research is needed to quantify the magnitude of this effect, as previous research (e.g., Kelly et al. 2021) has found that perfect knowledge of the field and climate conditions is not required for near-optimal irrigation scheduling.

To balance computational time and resources, our analysis focused on a selection of key factors (adaptation frequency, foresight of near-future weather, abstraction restrictions) that we hypothesised may be important determinants of the value of adaptive irrigation scheduling approaches. Nonetheless, other factors not analysed in the present paper may also influence the value of adaptive scheduling. For example, if the baseline fixed irrigation strategy is sub-optimal or poorly defined (Kelly et al. 2021), there may be greater value in updating this strategy throughout the season. The presence of multi-year quotas, rather than the single-year quotas considered in the present article, may also affect the value of adaptation as in-season decisions have implications for water availability in later years and thus make the optimization process both an intra- and inter-seasonal problem (Young et al., 2021). Furthermore, the presence of financial safety nets such as crop insurance (USDA, 2022; Suchato et al., 2022) may reduce the value of adaptation as outlier drought years, in which the largest benefits of adaptation are observed in this study, will result in profit losses that are less severe for farmers.

Finally, it is important to acknowledge that results presented in this study are dependent on the ability of the underlying crop-water model to accurately represent real-world crop responses to soil moisture deficits, weather conditions, and irrigation management practices. AquaCrop has been shown to adequately capture maize yield responses to water stress within our study area and the United States generally (Heng et al., 2009; Sandhu & Irmak, 2019). Further research should extend the analysis presented here to consider how the value of adaptive irrigation scheduling is influenced by the choice of underlying crop model and parameterization, including to consider potentially important differences that may exist across different crop varieties (e.g., more/less drought tolerant), soil types (e.g., higher/lower water holding capacity), or climatic regimes (e.g., greater/lesser rainfall variability).

4.5 Conclusions

The aim of this manuscript was to assess the potential added value of in-season adaptation of irrigation scheduling practices for improving productivity and profitability of agricultural water use under climate uncertainty. For a case study of maize production in Nebraska, our analysis demonstrates the vast majority of the potential profits can be achieved by adopting

a fixed irrigation strategy that is optimised pre-season to account for expected variability in weather outcomes but which is not adjusted within each year. Increases in average profit, as well as decreases in profit variance, are possible by re-optimizing this strategy within the season, with the size of this increase dependent on the potential frequency of adaptation and foresight of weather for the days ahead. However, the overall magnitude of these increased returns are small, even in the presence of restrictions on total seasonal water use by farmers. As a result, our findings suggest that farmers and water managers may be better to prioritise the adoption of simpler average irrigation scheduling heuristics that are well optimised to local production conditions and constraints. In doing so, farmers are likely to achieve the majority of the potential gains in crop yields and water use savings, while also minimizing costs associated with implementing more complex adaptive scheduling approaches and avoiding potential risks of maladaptive outcomes.

Acknowledgments and Data

The work contained in this article was funded by the National Environmental Research Council's Understanding the Earth, Atmosphere, and Ocean Doctoral Training Programme, Grant NE/L002469/1.

To support extensions of our research in line with future research directions outlined above, all code used in the analysis will be available in a public github repository (https://github.com/thomasdkelly/adaptive-irrigation), allowing anyone to duplicate the analysis.

SUPPORTING INFORMATION

This supporting information includes Figures S1 through S3 and Tables S1 and S2. Figure S1 provides insight into the choice of the number of optimization repetitions. Figure S2 shows how the regression relationships change with the seven-day perfect weekly forecast framework. Table S1 displays all the optimised *fixed* and *potential* thresholds. Table S2 displays all the optimised *fixed* thresholds for each irrigation cap.

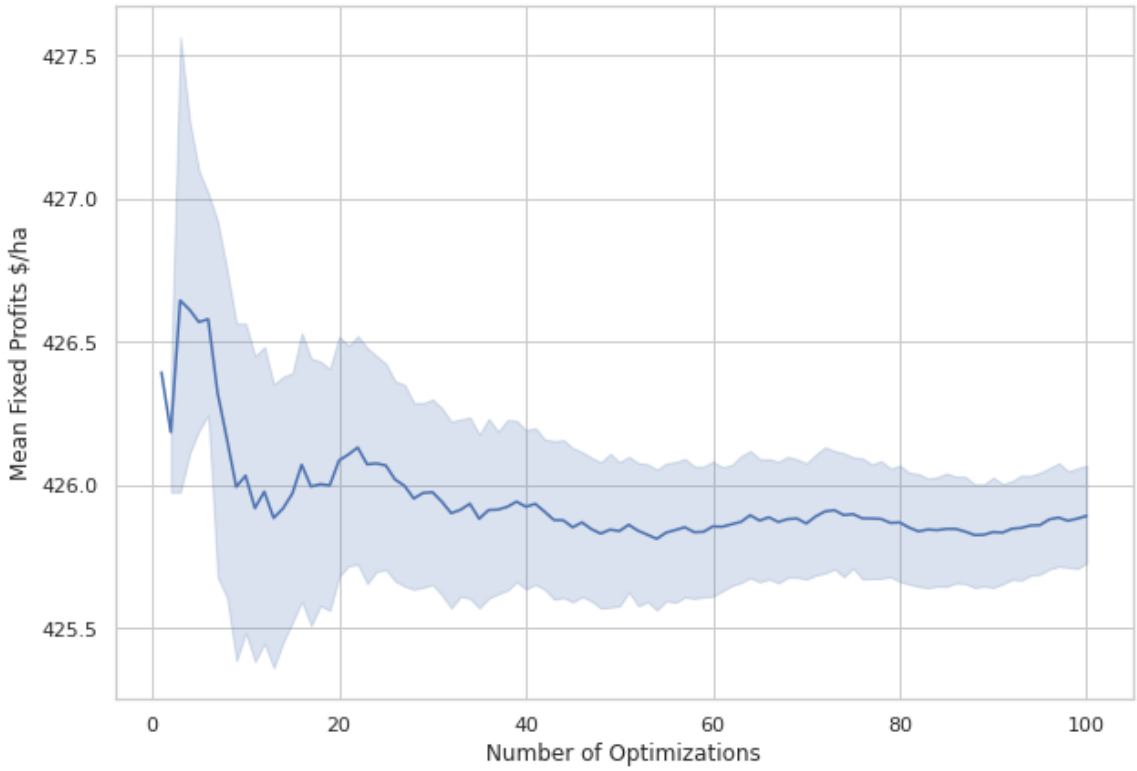

Figure 4.S1. The average profit over all repetitions achieved by the optimiser as a function of the number of repetitions. The shaded area represents the 95% confidence interval for the mean value across the number of repetitions specified.

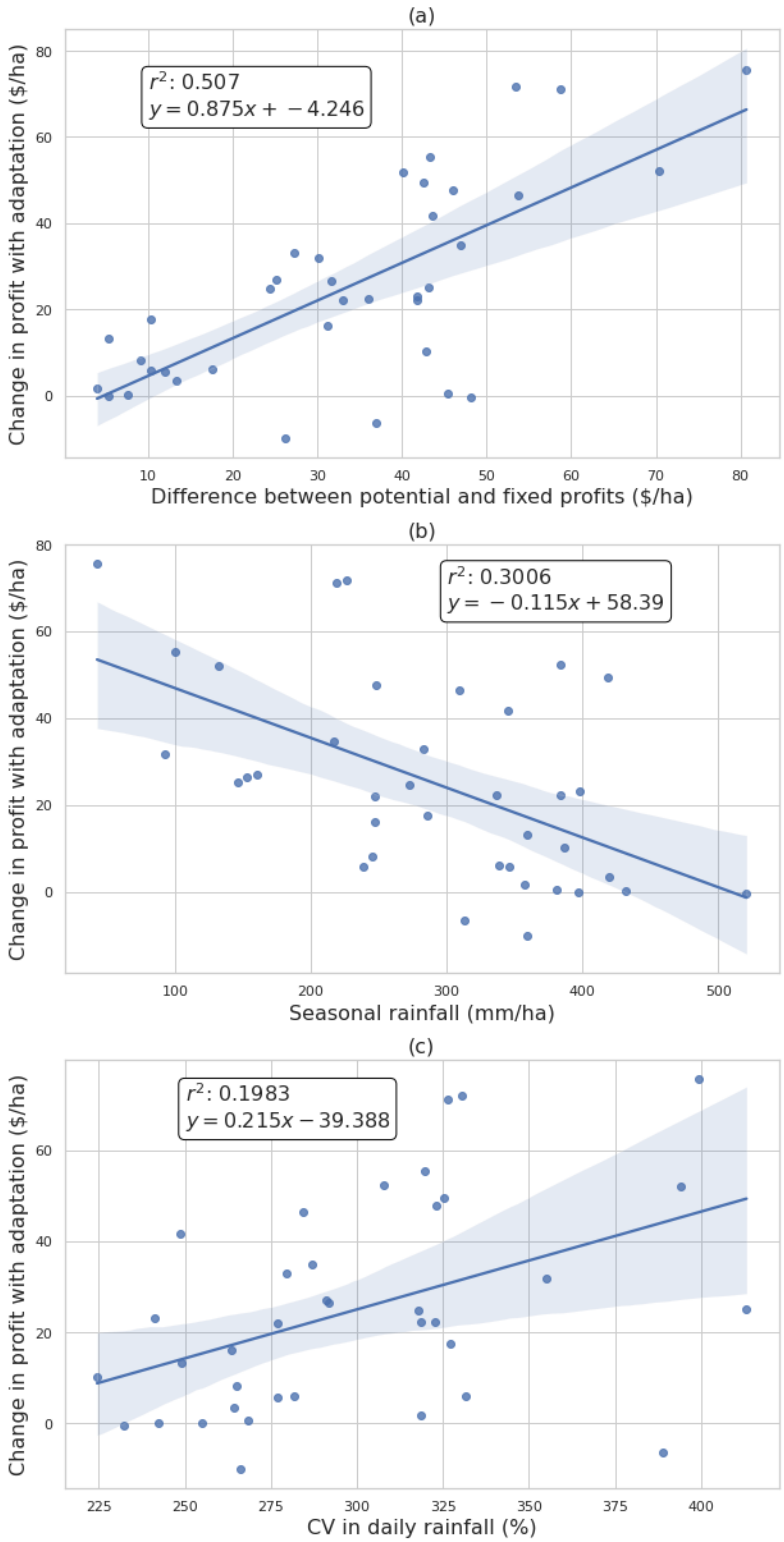

Figure 4.S2. The relationship between: (a) adaptation value and quality of starting strategy, (b) adaptation value and total seasonal rainfall, (c) adaptation value and CV in daily rainfall. Results are displayed for the 7-day adaptation strategy with perfect weekly forecast. Each point represents one of the 37 climate years, and adaptation value is averaged over 25 repetitions of the analysis framework. Shaded areas represent the 95% confidence interval for the linear regression.

Table 4.S1. *Fixed* and *potential* strategies (and resulting profits) for each climate year. Strategies consist of a set of soil-moisture thresholds.

Year	Rainfall (mm/ha)	SMT1 (%TAW)	SMT2 (%TAW)	SMT3 (%TAW)	Fixed profit ($/ha)	Potential profit ($/ha)
1982	397.42	43.42	83.38	32.54	587.83	593.22
1983	152.57	38.72	69.74	34.22	191	222.69
1984	92.92	44.93	87.31	36.5	209.49	239.58
1985	146.18	59.57	64.84	35.78	309.61	352.77
1986	381	26.81	61.43	32.11	470.54	515.92
1987	359.27	55.27	63.17	35.16	555.14	560.45
1988	345.38	41.9	70.94	38.59	366.29	409.93
1989	272.48	24.89	70.61	35.43	465.3	489.73
1990	217	40.69	67.82	40.73	277.73	324.64
1991	384	22.32	71.44	34.15	442.65	484.38
1992	419	57.05	18.94	43.73	755.7	798.21
1993	420	42.41	43.86	42.18	738.45	751.77
1994	219	62.58	50.27	35.82	347.52	406.19
1995	357	37.76	62.41	36.05	468.65	472.66
1996	520.77	31.61	19.48	21.42	696.82	744.98
1997	248	22.34	56.11	34.05	373.14	419.25
1998	283.36	52.09	64.45	55	356.07	383.37
1999	345.78	45.62	57.6	30.6	466.47	476.79
2000	100	35.74	62.9	34.03	192.31	235.58
2001	245	55.65	72.99	34.94	312.96	322.1
2002	132.49	70.51	75.34	35.47	164.75	204.81
2003	160	54.37	59.09	35.82	241.93	267.12
2004	397.88	68.78	62.01	33.36	640.19	681.93
2005	312.63	35.87	65.9	51.66	439.86	476.75
2006	226.07	61.33	51.04	30.36	357.79	411.3
2007	286.23	50.15	55.59	33.86	400.82	411.1
2008	336.29	36.17	72.59	34.79	567.55	603.54
2009	431.56	50.66	69.74	37.11	699.33	706.98
2010	309.57	56.21	50.74	37.59	410.89	464.68
2011	359.62	13.75	63.34	37.19	419.88	446.13
2012	42.65	27.26	77.68	39.12	35.18	115.88

2013	239.01	21.93	60.81	35.1	364.37	376.33
2014	386.33	27.17	75.4	35.56	589.6	632.38
2015	384.24	43.29	56.83	40.83	459.13	529.49
2016	247.43	26.33	84.75	33.92	428.88	460.15
2017	247.32	26.49	68.25	40.93	417.45	450.45
2018	338.46	2.78	62.85	36.33	539.81	557.48
Fixed		50.47	61.43	36.22	425.975135	

Table 4.S2. Optimal fixed strategies for different levels of irrigation cap.

Irrigation Cap (mm/ha)	SMT1 (%TAW)	SMT2 (%TAW)	SMT3 (%TAW)	Fixed Profit ($/ha)
75	12.52	16.89	33.18	-347.41
100	13.51	7.85	81.46	-215.97
125	5.13	3.72	80.04	-101.17
150	3.13	27.3	70.47	-5.55
175	1.88	33.51	78.61	67.4
200	40.57	18.27	79.37	133.69
250	49.35	44.28	34.53	261.99
300	49.57	40.31	37.51	339.08
350	46.9	44.28	37.42	374.58
400	50.56	48.61	39.18	395.6
450	50.46	61.64	36.22	413.04
500	50.41	61.62	36.22	418.16
550	50.46	61.45	36.22	419.62
No cap	50.47	61.43	36.22	

5 ASSESSING THE VALUE OF DEEP REINFORCEMENT LEARNING FOR IRRIGATION SCHEDULING

This Chapter presents the article: *Assessing the value of deep reinforcement learning for irrigation scheduling.* This article is targeted for submission to *Computers and Electronics in Agriculture* in Summer 2022.

My contribution to the article was as follows: (1) Development of the research questions and methodology. (2) All programming required to run the framework and analyse results. (3) Analysing results and producing tables and figures. (4) Writing the contents of the article including literature reviews.

Assessing the value of Deep Reinforcement Learning for irrigation scheduling

T. D. Kelly[1], T. Foster[1], David M. Schultz[2, 3]

[1] Department of Mechanical, Aerospace & Civil Engineering, University of Manchester, Manchester, UK.

[2] Department of Earth and Environmental Sciences, University of Manchester, Manchester, UK.

[3] Centre for Crisis Studies and Mitigation, University of Manchester, Manchester, UK.

Corresponding author: Thomas Kelly (Thomas.kelly-2@manchester.ac.uk)

ABSTRACT

Due to increasing global water scarcity pressure, researchers, policy makers and industry are looking for innovative solutions to increasing agricultural water productivity. Motivated by recent success within complex decision making environments, Deep Reinforcement Learning (DRL) is being proposed as a method for optimizing irrigation strategies. Early research has hinted towards increased profits with DRL compared to heuristic approaches such as single soil-moisture thresholds or fixed schedules. However, an assessment of the value of DRL for irrigation scheduling that incorporates local climate variability and water-use restrictions has yet to be performed. To address this gap in the literature, we created *aquacrop-gym*, an open-source Python framework for researchers to train and evaluate customised irrigation strategies within the crop-water model AquaCrop-OSPy. In this analysis, *aquacrop-gym* was used to simulate irrigated maize production in the central United States. The DRL and heuristic approaches were both trained on 70 years of weather data produced from the weather generator LARS-WG, and evaluated on 30 unseen validation years of generated weather data. Findings from this analysis show that in the presence of high rainfall variability, DRL does not outperform the optimised heuristic. However, in the scenario where rainfall is set to zero, DRL approaches achieve higher profits on the unseen validation years. Similarly, DRL approaches also outperform optimised heuristics when severe water-use restrictions are introduced. Our analysis demonstrates that DRL approaches are a promising method of irrigation scheduling, but have not yet been shown to be an overwhelming improvement compared with optimised heuristics.

5.1 INTRODUCTION

Agriculture is the main consumptive user of freshwater in many water-scarce regions, and thus a sector that is a both a cause and a victim of water scarcity. Improving agricultural water productivity is therefore an essential component of efforts to address water scarcity worldwide, as well as meeting food insecurity challenges (Howell, 2001; Fereres & Soriano, 2006).

Efforts to improve agricultural water productivity are fundamentally focused on tailoring irrigation to the crop-water needs as well as minimizing non-beneficial losses of water. Effective scheduling of irrigation during the growing season, alongside other interventions

such as improved irrigation technologies or land management practices, are critical to achieving improvements in agricultural water productivity. Commonly, farmers schedule irrigation based on rules or heuristics, such as soil-moisture levels or crop conditions (USDA-NASS, 2018; Blonquist et al., 2006; Linker & Sylaios, 2016; Foster & Brozović, 2018; Kelly et al., 2021), which may be constant throughout the season or vary depending on crop growth stage. However, while irrigation strategies based on simple rules and heuristics are an intuitive and agronomically logically way to schedule irrigation, they may lead to sub-optimal irrigation decisions if the underlying assumptions on which the heuristic is designed change or vary unexpectedly. For example, an unexpected rainstorm or longer than usual period of drought may render the current heuristic sub-optimal in a given year. Similarly, an unanticipated change in water availability (e.g., due to equipment failure, emergency drought measures) may render a farmers irrigation heuristics sub-optimal given altered water supply risks.

One solution to this challenge is to adapt irrigation scheduling heuristics by re-optimizing decision rules and triggers at various points within the season. Research has demonstrated that combining crop simulation models with field monitoring, weather forecast, and other data can be used to support adaptation of irrigation scheduling and improve yields, water productivity, and profitability (Wang & Cai, 2009b; Cai et al., 2011; Jamal et al., 2019; Linker, 2021). However, this approach adds additional computational overhead during the season, especially if re-optimizing frequently throughout the season, thereby increasing costs for producers or irrigation scheduling service providers. In the case of service providers, re-optimizing strategies for tens of thousands of different cropping systems in near real-time may not be economically or technically feasible. Furthermore, as well as these issues with in-season adaptation, incorporating the increasing amounts of data available to farmers into heuristics will require large numbers of complex, interconnected rules that will be difficult to design and optimise.

Addressing these issues, artificial intelligence based approaches such as Deep Reinforcement Learning (DRL) (Sutton et al., 1998) has been suggested as way to automatically and flexibly adjust and adapt decisions during the season in response to any relevant information. DRL approaches have successfully surpassed the best human benchmarks in other complex decision making environments such as video and board games

(Silver et al., 2018; Vinyals et al., 2019; OpenAI et al., 2019; Badia et al., 2020). Recent research has explored the viability of using DRL for irrigation scheduling (Yang et al., 2020; Chen et al., 2021; Alibabaei et al., 2022). In comparison with non-optimised schedules and heuristics, DRL methods were found to deliver improvements in water productivity and profits, highlighting their potential value as a tool for real-world irrigation management. However, in all of these works the DRL agent was evaluated on a limited number of unseen test seasons (less than 3 for all of these analyses), which is insufficient to assess performance across diverse and uncertain weather scenarios faced by farmers in real-world production environments. Additionally, trained DRL agents were compared only to handpicked schedules (i.e., a set of dates and depths chosen by researchers), fixed time intervals (i.e., irrigate 25 mm every *N* days), or single value heuristics (i.e., irrigate if the soil-water content drops a specified value), where this value was chosen by researchers rather than optimised for the specific production environment. These choices of benchmark irrigation strategy limit our understanding of the added value of using DRL compared to the optimised heuristics commonly used in irrigation scheduling research (Linker et al., 2016; Linker & Kisekka, 2017; Kelly & Foster, 2021).

In this article, we conduct a more systematic assessment of the value of DRL approaches to irrigation scheduling in comparison to conventional optimised heuristics, across a range of climate conditions. We assess how design choices, such as when and how often the DRL agent makes irrigation decisions, as well as assumptions about water use restrictions, impacts estimates of productivity and profitability gains from DRL. Through these analyses, we provide new insights into how best to evaluate the benefits of DRL for irrigation scheduling, as well as discussions about the conditions under which DRL may be a powerful tool in improving irrigation water productivity.

5.2 METHOD

In this section we describe the framework used to train and evaluate DRL agents for irrigation scheduling. First, we give a brief overview of DRL (Section 5.2.1), then describe how we apply DRL to irrigation scheduling within a crop-simulation model (Section 5.2.2). The chosen case study, as well as experimental details are then described (Section 5.2.3), followed by a discussion on overfitting and how to report DRL performance (Section 5.2.4).

Finally, we describe the optimised heuristic that we will use to assess the added value of DRL (Section 5.2.5).

5.2.1 Deep Reinforcement Learning

A typical reinforcement learning problem is framed as the interaction between an agent and an environment (Figure 5.1). At each timestep, the agent receives some observation of the environment *O*, and makes action *A* in response. This action has some impact on the environment which then returns a reward *R* along with the next observation *O*. This loop continues until a terminal state of the environment is reached, after which point the *episode* has finished. The goal of the reinforcement learning agent is to learn a mapping of observation to action that maximises the sum of future rewards (Sutton et al., 1998).

In Deep Reinforcement Learning (DRL), the mapping of observation to action is an artificial neural network, with the weights of that neural network being continually updated in response to the rewards received by the agent. In this analysis we use Proximal Policy Optimization (PPO) (Schulman et al., 2017), which is one of the most popular DRL algorithms and has been used in a wide variety of environments such as robotics (Andrychowicz et al., 2019), video games (OpenAI et al., 2019) and flow control (Rabault et al., 2019). PPO also has the advantage that it can be easily used for both continuous and discreet action spaces, allowing us to define many different types of irrigation strategies. Early experimentation found PPO provided the fastest and most stable training performance for this analysis compared to other algorithms tested. The code implementation of PPO that we use is from the DRL Python library *ray-rllib* (Liang et al., 2017). This library is an actively maintained, production-ready framework that is currently used within several industries (https://docs.ray.io/en/latest/rllib/index.html).

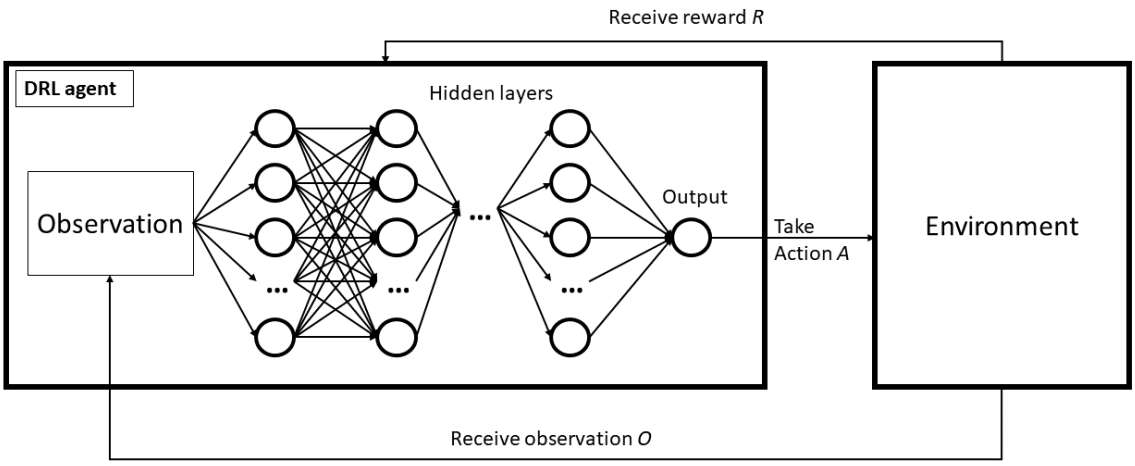

Figure 5.1. Typical framing for a reinforcement learning problem, where a DRL agent aims to maximise the sum of future rewards.

5.2.2 Applying DRL to irrigation scheduling

To train DRL agents to make irrigation decisions, a crop-simulation environment is needed to rapidly test different irrigation strategies. Moreover, this environment needs to follow the *Observation, Action, Reward, next Observation* (*O–A–R–O*) framework outlined in the previous section. For this analysis, we use AquaCrop-OSPy (Kelly & Foster, 2021) to represent the cropping system, and serve as the environment with which the DRL agent can interact with, make irrigation decisions, and receive feedback on how these decisions impact seasonal profits. AquaCrop-OSPy is an open source Python implementation of the crop-water productivity model AquaCrop (Steduto et al., 2009). Developed by the UN-FAO, AquaCrop is one of the most widely used crop simulation models available, featuring in over 334 peer reviewed publications in the ten years since its launch (Salman et al., 2021). AquaCrop aims to accurately represent crop response to water stress, whilst also reducing complexity and input data requirements, enabling it to be used in data scarce environments. AquaCrop-OSPy has been designed to allow flexible and complex irrigation decisions to be made outside of the model, allowing straightforward integration with DRL Python libraries or optimization algorithms. To facilitate this interaction between DRL agent and AquaCrop-OSPy, we created the *aquacrop-gym* open-source Python library (https://github.com/aquacropos/aquacrop-gym).

To train a DRL agent within *aquacrop-gym*, a DRL agent is first initiated with random neural network weights. To gather information with which to learn from, the agent then performs

a number of *training rollouts* (Figure 5.2). For each *training rollout* a weather year is randomly sampled from the training set and an AquaCrop-OSPy simulation is initialised. Every *N* number of days within the simulation, AquaCrop-OSPy will pass to the agent a set of observations (e.g., current soil-moisture content). These observations are then processed by the agent and an action is recommended (e.g., irrigation depth). The simulation then runs for another *N* days after which the agent will output another action based on the latest observation. At the end of the season, the final reward (e.g., seasonal profits) is calculated and sent back to the agent.

After the *training rollout* is complete, all of the saved (*O*–*A*–*R*–*O*) transitions are used to update the weights of the agent's internal neural network. Once the agent's strategy has been updated, a *validation rollout* (Figure 5.2) is performed. The *validation rollout* evaluates the agent against all of the unseen validation years and calculates the mean profit achieved by the agent. Since none of the (*O*–*A*–*R*–*O*) transitions are saved during this *validation rollout*, no learning occurs and only the mean reward is saved to track the learning progress of the agent. This loop of: collect training data, update strategy, validate strategy will continue for a specified number of iterations or until manually stopped.

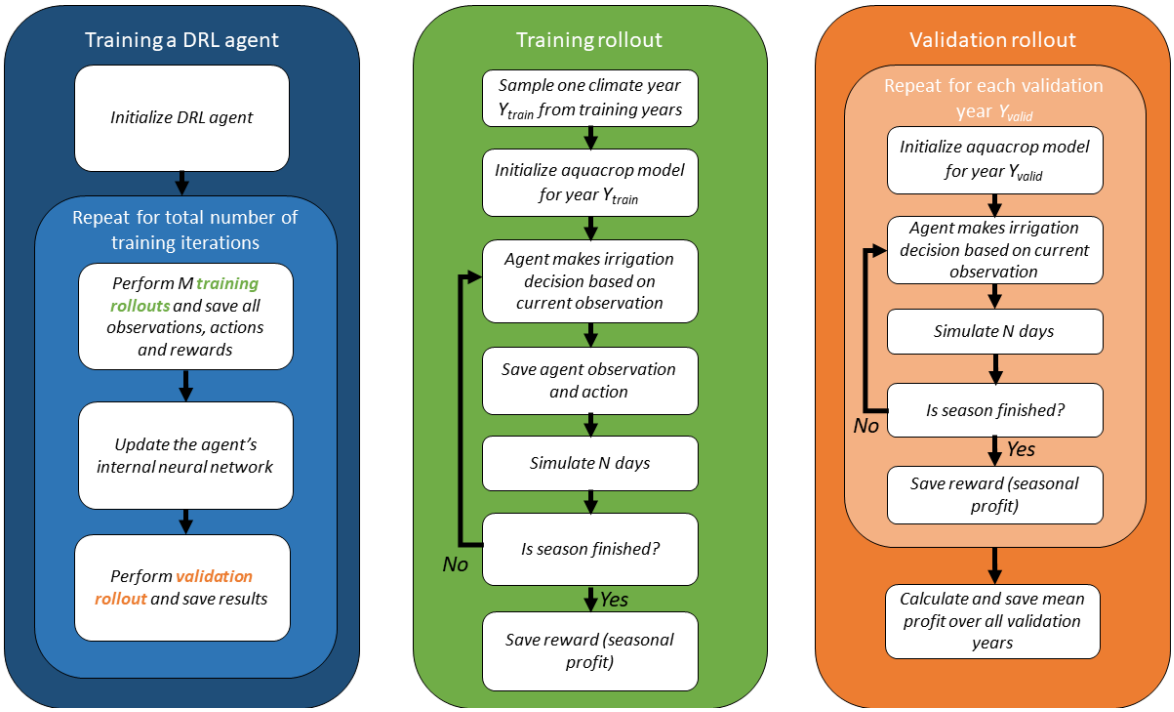

Figure 5.2. Flow diagram detailing how training data is collected by the DRL agent, then used to update the agent's decision making. After updating the new strategy is evaluated on the validation set.

5.2.3 Case study, DRL design choices and experimental details

For the case study for this analysis, we use AquaCrop-OSPy to simulate irrigated maize production in southwest Nebraska. Maize is the most commonly grown and irrigated crop within Nebraska, which itself contains more irrigated acres than any other state in the United States (USDA-NASS, 2018). Irrigation water in Nebraska is primarily sourced from the High Plains aquifer, which has experienced severe declines in aquifer storage as a result of over abstraction (Scanlon et al., 2012; McGuire, 2017). In response to water scarcity issues, water policy makers within Nebraska have imposed water allocations on producers, to limit their abstractions for irrigation (Young et al., 2021).

To train and validate our model, we utilise historic weather data for southwest Nebraska from a monitoring station in Champion, Nebraska (HPRCC, 2016). This weather station provides 37 years (1982-2018) of observational data (daily minimum and maximum temperature, precipitation, and solar radiation), which we subsequently input to the

weather generator LARS-WG (Semenov & Barrow, 1997) to generate a longer 100-year synthetic time series for our analysis. This weather generator determines the statistical characteristics of a given daily weather time series and uses these characteristics to generate new time series of daily weather data. LARS-WG has been shown in multiple recent studies to successfully represent the weather characteristics (e.g., monthly mean and variance of temperature or precipitation) of the input time-series (Agarwal et al., 2014; Awal et al., 2016). Using this method, we greatly increase the number of weather years for training (70 years) and validating (30 years) whilst still adequately representing the observed historical weather conditions (Figure 5.S1). Other inputs, including crop and soil parameters, for AquaCrop simulations are drawn from the literature, specifically from previous studies that use the model to simulate maize production in the central United States (Foster et al., 2015; Linker & Kisekka, 2017; Sandhu & Irmak, 2019), For calculation of seasonal profits at the end of each simulation (Equation 4.1), we assume a crop price of $180 per tonne (USDA, 2019), irrigation costs of $1 per ha-mm (CropWatch, 2019a), and fixed non-irrigation production costs of $1728 ha^{-1} (CropWatch, 2019a). Main DRL hyperparamaters are specified in Table 5.S1, and full details of source code, model and DRL parameters are all available at (https://github.com/aquacropos/aquacrop-gym).

Determining what observations the agent should receive, which actions it is allowed to take, and what reward it will maximise are critical design choices that will depend on the cropping system being simulated. Examining all possible combinations of observation and action sets is out of the scope of this analysis. However, our aim is to test several of the most common and likely combinations that are applicable for our study area in the central United States. The range of action sets chosen for our analysis are displayed in Table 5.1. For *depth discreet*, the agent choses an irrigation depth to apply on the following day from a set of choices (0, 2.5, 5, …, 22.5, 25 mm ha^{-1}). For *depth continuous*, the agent can chose any depth to apply the following day within the specified bounds [0, 25 mm ha^{-1}]. For *binary*, the agent decides whether or not to apply 25 mm of irrigation the following day. For *soil-moisture threshold*, the agent decides the soil-moisture threshold to use until the next decision step. If the soil-moisture content ever drops below this threshold, irrigation is triggered to fill the soil back to field capacity (subject to the maximum irrigation application of 25 mm ha^{-1}).

Finally, for all of these action sets, we compare four different decision intervals (i.e., the number of days between actions) 1, 3, 5 and 7 days.

In each case, observations provided to the agent before each irrigation decision are as follows: mean daily reference evapotranspiration and precipitation for the season so far; and for the past 7 days only, current root-zone soil-moisture content, canopy cover, biomass growth, cumulative irrigation, water supply remaining, crop growth stage, and the total number of days since the season began. These variables are either obtained through the input weather time series or outputted by the AquaCrop-OSPy model. This choice of variables represents a range of weather, soil, and crop condition information that could plausibly be monitored or collected by farmers to support a DRL based scheduling approach, and could be extended or reduced depending on the monitoring capabilities of the individual production system being modelled.

Table 5.1 The action sets this case study compares.

Action set	Bounds	Example action
discreet depth	[0,25] mm ha^{-1}	0, 2.5, … 22.5, 25 mm ha^{-1}
continuous depth	[0,25] mm ha^{-1}	8.43 mm ha^{-1}
binary	[0,25] mm ha^{-1}	0, 25 mm ha^{-1}
soil moisture threshold	[0,100] % TAW	63.2%TAW

As discussed in the previous section, DRL agents will be trained on 70 years of training data, and evaluated on 30 years of validation data. The mean and standard deviation in seasonal rainfall (May–September) for the training and validation sets were closely comparable, equal to 309.6±93.7 and 306±93.2 mm ha^{-1} respectively. Although both the training and validation sets were drawn from the same statistical distribution, the individual time series will be different, and thus provides a way for us to assess how well the agent might perform if deployed in the real world under uncertain weather conditions. For a similar case study of irrigated maize production in Nebraska, Chapter 4 of this thesis found the rainfall variability in this region resulted in large heterogeneity in seasonal profits. To assess the impact of this local variability on DRL performance, an additional experiment was therefore conducted

where the rainfall was removed completely from the training and validation sets. We hypothesise that by removing this important source of weather variability, and hence reduce the difference between the training and validation time series, the performance of DRL should improve compared to when rainfall is present.

Previous assessments of DRL for irrigation scheduling have assumed unrestricted irrigation, whereby the DRL agent can irrigate as much as required. However, this assumption is not realistic to our case study region of Nebraska, or even other regions globally where water use may be restricted due to regulatory policies such as yearly or multi-yearly quotas and allocations (Berbel et al., 2019; Young et al., 2021). Given that the marginal gains from applying additional water are greatest in times of water scarcity, we hypothesise that more complex strategies such as DRL may be more valuable when water resources are limited. To assess the impact of water use restrictions on the value of DRL, the framework presented so far is repeated and water use limits of 75, 100, 125 mm ha^{-1} are imposed on all simulation years. This restriction will prevent any further irrigation once the cap has been reached in any given growing season.

To summarise, four DRL agents will be trained initially to compare the action sets described in Table 5.1. The best performing action set will be used to train an additional four DRL agents to compare the four decision intervals 1, 3, 5 and 7 days. With this chosen DRL configuration (i.e., action set and decision interval), two more experiments will be conducted – one with no change to local rainfall and one with rainfall set to zero. Finally, with the same DRL configuration, three experiments will be conducted for each seasonal water use limit 75, 100, 125 mm ha^{-1}. In total this will result in 13 experiments.

5.2.4 Learning curves, evaluation criteria and overfitting

One of the downsides of using an artificial neural network based irrigation strategy is that neural networks, if trained for long enough, will often overfit to training data. When overfitting occurs, the training set performance continues to increase but the performance on the unseen validation set begins to decrease. This divergence occurs as the neural network begins to find patterns that apply only to the training set, rather than general patterns that apply also to the validation set. As a result of this overfitting, it is often desirable to stop training before reaching an optimum score on the training set – by which

point the model has very likely overfit. To track learning process, throughout training the DRL agent will be continually evaluated on both the training and validation years, allowing us to monitor when and if overfitting occurs.

Figure 5.3 displays these evaluation results for a DRL agent that makes *discreet depth* decisions every 7 days. Figure 5.3a is generally referred to as the learning curve and tracks the agent's performance as a function of training steps (i.e., number of irrigation decisions made and learned from). Performance on the training set continues to increase throughout. However, the validation set performance declines after approximately 500,000 training steps. The divergence between training and validation performance highlights an important decision that needs to be made regarding how to report the final DRL performance, with which to compare against other irrigation strategies. Once choice is to report the highest validation profits achieved by the agent, which would give an indication of the potential of DRL if training is stopped at the right moment. However, this choice would overestimate the performance of DRL as it assume knowledge of the validation years, which would not be available to a farmer as the validation set represents future unknown weather. Another choice is to use the final validation set profit after training has stopped, which would not assume any knowledge of the validation years. However, as shown in Figure 5.3a, there can be large variation in validation profits between learning updates, and so choosing just one evaluation result may be misleading. Therefore, we instead calculate the rolling mean across the previous 10 validation set evaluations (Figure 5.3b). When comparing the performance of DRL on the validation set, we will report the highest 10-evaluation average.

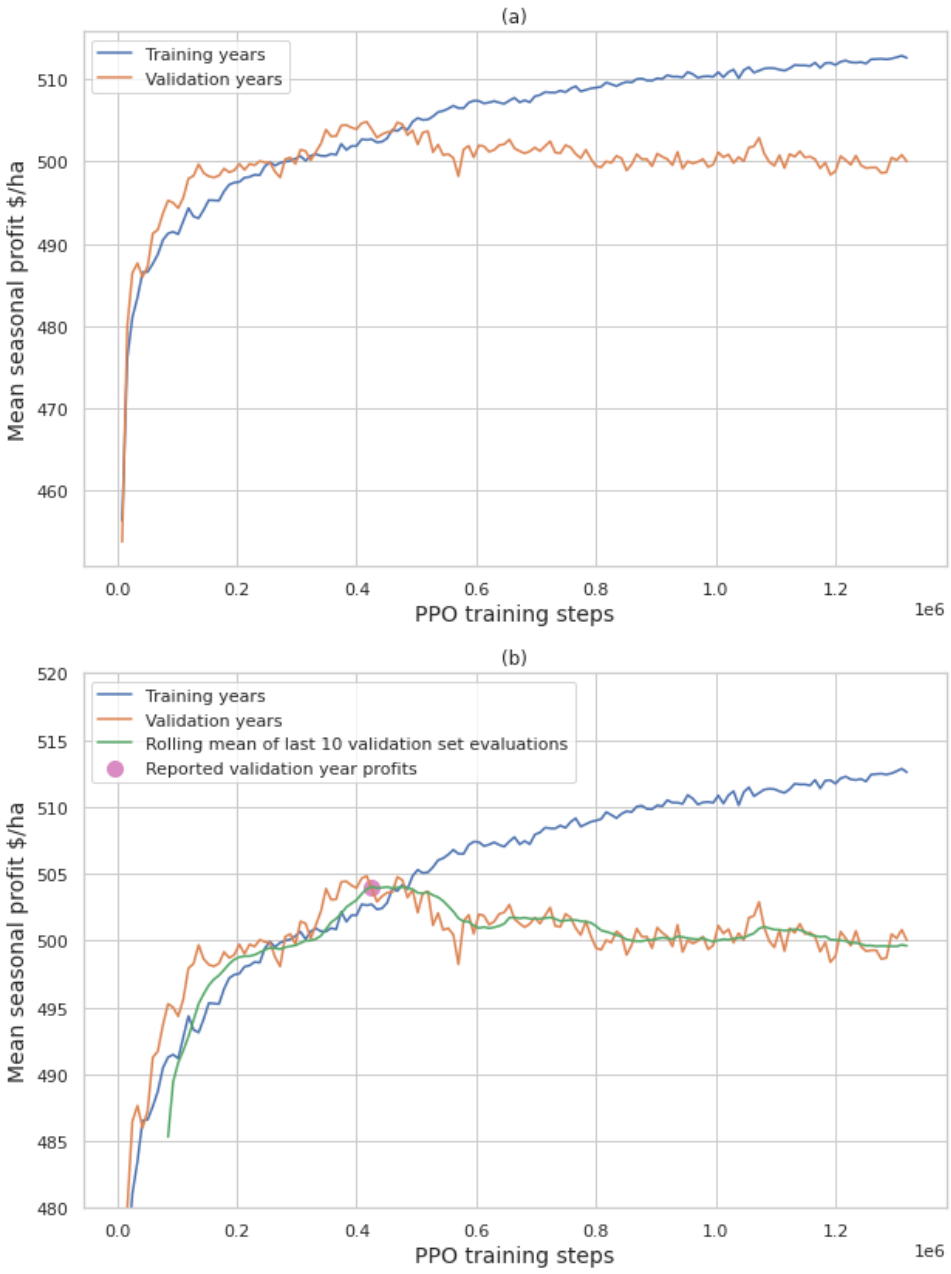

Figure 5.3 (a) Average training profits and testing profits as a function of training steps. (b) Average training profits and testing profits as a function of training steps, as well as a rolling average of validation profits over 10 evaluations. The highest point on this rolling average calculation will be reported as the validation profit.

5.2.5 Optimizing heuristic benchmark

Previous assessments of DRL for irrigation scheduling have only compared DRL agents with handpicked schedules or one-value non-optimised heuristics (Yang et al., 2020; Chen et al., 2021; Alibabaei et al., 2022). Chapter 4 of this thesis found that fixed optimised heuristics

trained on historical weather data can achieve over 90% of the potential profits (i.e., those achieved with perfect seasonal weather foresight) for a similar production system to the central United States case study used in the analysis presented in this chapter. The irrigation heuristic used in analyses in Chapter 4 was composed of a set of four soil-moisture thresholds, corresponding to the four main crop growth stages within AquaCrop. If the daily soil-water content drops below the current threshold at any point, irrigation is triggered. Within *aquacrop-gym* it is also possible to define this same four-soil-moisture threshold strategy and link this to an optimization algorithm. We use this approach to define a set of four soil-moisture thresholds that maximise average profit over the 70 training years. This set of thresholds will then be evaluated against the unseen 30 validation years and thus provide the benchmark against which to assess the added value of DRL. This benchmark achieves \$513 ha^{-1} on the training years and \$521 ha^{-1} on the validation years. The benchmark also achieves \$503 ha^{-1} when evaluated on the 37 years of historical data that was originally imputed into LARS-WG. By comparing these profits to those that could be achieved by directly optimizing the heuristic on those 37 historical years (\$506 ha^{-1}), we can be confident that our approach of using generated data to train and optimise strategies – as well as the choice of weather generator – produces robust strategies for the case study region.

Optimising the heuristic benchmark as well as running DRL training experiments were all conducted on a single AMD Ryzen 7 3700X, 3.60GHz, 8 core computer, with a GeForce RTX 3070 GPU. Training 1 million DRL training steps on this machine took approximately 4 hours, whereas optimizing a set of four soil-moisture thresholds took approximately 20 minutes.

5.3 RESULTS

This section begins by assessing how the performance of DRL varies depending on choices of action sets and decision frequencies (Section 5.3.1). Next, the effect of climate variability on DRL performance will be assessed (Section 5.3.2), followed by the introduction of restrictive caps on irrigation water use (Sections 5.3.3).

5.3.1 Comparing DRL configurations

Table 5.2 shows the mean seasonal profits achieved by different irrigation scheduling methods after 1 million training steps. DRL approaches, independent of these choices for

how irrigation could be triggered, achieved higher profits on the training set than the fixed optimised irrigation scheduling heuristic. However, the fixed heuristic achieved marginally higher profits than the DRL approaches over the validation set. This result illustrates that deep learning algorithms are able to learn complex irrigation strategies (as evidenced by the superior training set performance), but may not be able to transfer those superior strategies into improved performance in out-of-sample production scenarios (i.e., the validation set). The differences in profit between all strategies is extremely marginal and so these results alone cannot be used to definitively rule out any approach as a viable method of scheduling irrigation. The *discreet depth* action set achieved the highest validation profit of all DRL scheduling strategies, and so will be used as the basis for comparison with fixed soil-moisture threshold (SMT) heuristics for the rest of the analysis presented.

Table 5.2 Reported seasonal profits (highest mean profits over 10 evaluations) for different action sets.

Scheduling method	Training set profits $ ha^{-1}	Validation set profits $ ha^{-1}
DRL – *discreet depth*	518.9	520.1
DRL – *continuous depth*	**520.2**	512.6
DRL – *binary*	515.3	517.8
DRL – *soil-moisture threshold*	517.5	515.7
Fixed Optimised SMT	513.1	**521.4**

The second DRL design choice we examined was the decision interval (i.e., how often the DRL agent makes irrigation decisions). For the best performing action set (*discreet depth*), we train the DRL agent with decision intervals (1, 3, 5, 7 days) for 1 million training steps and compare the validation set performance. This comparison showed that a decision interval of 3 days resulted in the highest profits over both the training and validation set (Table 5.3). The variation in profits between different decision intervals was greater than the profit variations found between different action sets (Table 5.2). This result could indicate that how often the DRL agent makes decisions is more important than how the depth is specified, at least for our specific crop model environment and case study.

Table 5.3 Reported seasonal profits (highest mean profits over 10 evaluations) for different decision intervals using the *discreet depth* action set defined in Section 5.2.3.

Decision Interval	Training set profits $ ha^{-1}	Validation set profits $ ha^{-1}
1 day	492.2	497.6
3 days	**518.9**	**520.1**
5 days	515.1	511.0
7 days	510.4	504.0

5.3.2 Climate variability

Despite the range and distribution of weather scenarios in model training and validation being comparable, the results presented in Section 5.3.1 show that complex DRL scheduling algorithms have thus far only been able outperform simpler optimised soil-moisture thresholds on the training set. One reason for this underperformance is that specific weather scenarios encountered in the validation period are inherently different (e.g., in terms of specific sequencing of rainfall events), despite still having same overall statistical distribution (e.g., in terms of total seasonal rainfall). By setting the rainfall to zero, and hence reduce the difference between the training and validation time series, we expect the performance of DRL to improve on the validation set. To test this hypothesis, we removed the rainfall from both the training and validation weather dataset used within AquaCrop-OSPy, and retrained the DRL agent as well as re-optimised the fixed heuristic benchmark.

Table 5.4 shows the difference in profits between the DRL agent and the optimised soil-moisture thresholds when rainfall is included or removed. When rainfall was removed, optimised fixed soil-moisture thresholds achieved a profit of $300.6 ha^{-1} on the validation set, with reduced profit relative to that found in Section 5.3.1 due to larger irrigation amounts needed to meet crop water requirements across all years. In contrast, the DRL agent taking *discreet depth* actions every 3 days achieved $307 ha^{-1}. This result indicates that removing the biggest source of climate variability altered the added value of DRL scheduling relative to fixed SMT heuristics from a reduction of $1.27 ha^{-1} to an increase of $6.38 ha^{-1}, the latter a 2.1% increase in profits compared to the optimised heuristic.

Table 5.4 Increase in seasonal profits compared to optimised soil-moisture thresholds with and without rainfall.

Weather data	Increase in training set profits $ ha^{-1}	Increase in validation set profits $ ha^{-1}	Percentage increase in validation set profits
With rainfall	5.87	-1.27	– 0.2%
Without rainfall	4.07	6.38	+ 2.1%

5.3.3 Water use restrictions

The DRL and heuristic strategies evaluated so far have operated under unrestricted irrigation, whereby the DRL agent can irrigate as much as required. However, this assumption is not realistic to our case study region of Nebraska, which imposes yearly or multi-year restrictions on irrigation water abstractions. Given that the marginal gains from applying additional water are greatest in times of water scarcity, we hypothesise that more complex strategies such as DRL may be more valuable when water resources are limited. To assess the impact of water use restrictions on the value of DRL, the framework presented so far was repeated and water use limits of 75, 100, 125 mm ha^{-1} were imposed on all simulation years. This restriction will prevent any further irrigation once the cap has been reached in any given growing season. Table 5.5 shows the difference in profits between the DRL agent and the optimised soil-moisture thresholds for these different water use restrictions. Despite only small profit increase, DRL does provided added value in these water restricted scenarios compared to optimised heuristics. By introducing the most restrictive cap evaluated (75 mm ha^{-1}), the value of DRL increased from -$1.3 ha^{-1} to +$12.2 ha^{-1}, which equates to a 5.6% increase in profits compared to the optimised heuristic.

Table 5.5 The effect of water use restrictions on validation set profits and added value of DRL compared to optimised heuristics.

Irrigation cap mm ha^{-1}	DRL validation set profits \$ ha^{-1}	Increase in validation set profits \$ ha^{-1}	Percentage Increase in validation set profits
75	228	12.2	5.6%
100	286.5	7.1	2.5%
125	332	4.3	1.3%
No cap	520.1	-1.3	– 0.2%

5.4 Discussion

Despite the success of optimised heuristics within irrigation scheduling, incorporating the increasing amounts of data available to farmers into heuristics will require large numbers of complex, interconnected rules that will be difficult to design and optimise. Addressing these issues, artificial intelligence based approaches such as Deep Reinforcement Learning (DRL) (Sutton et al., 1998) has been suggested as way to automatically and flexibly adjust and adapt decisions during the season in response to any relevant information. Recent publications have found Deep Reinforcement Learning (DRL) can be used to learn irrigation strategies within crop-simulation frameworks (Yang et al., 2020; Chen et al., 2021; Alibabaei et al., 2022).

In this article we assessed the added value of using DRL for irrigation scheduling in comparison to optimised soil-moisture thresholds. The first part of the analysis found that allowing the DRL agent to choose between a range of discreet irrigation depths (i.e., 0, 2.5, … 22.5, 25 mm) lead to the best performance over the validation years. Despite performing the best compared to other DRL design choices, this was still not enough to surpass the validation profits achieved by the optimised soil-moisture thresholds. When rainfall was removed from the weather data, thereby decreasing the difference in weather time series between training and validation sets, the DRL agent was now able to increase validation year profits by 2.1% compared to optimised heuristics. Finally, when restrictive irrigation

caps were introduced (e.g., 75 mm ha^{-1}), the DRL agent was able to increase validation year profits by 5.6% compared to the optimised heuristics.

Our analysis found that under normal climate variability for the case study region and unrestricted water use, optimised soil-moisture thresholds achieved a higher average profit over the validation years than the DRL agent. This result broadly agrees with Alibabaei et al., (2022), who found that even a non-optimised fixed heuristic achieved higher profits than their DRL agent during one of their two test years. This marginal difference in validation year profits found between DRL and optimised heuristic (– \$1.3 ha^{-1}) cannot be taken as an exact quantitative comparison on its own, as another set of validation years could cause this value to vary slightly. Differences between how DRL performance is reported – i.e., when training is stopped, how validation set profits are calculated, and how many training repetitions are performed – can also lead to this difference in profits to vary slightly. However, the qualitative result that DRL agents struggle to surpass optimised thresholds is likely to remain, and was largely neglected in previous DRL assessments (Yang et al., 2020; Chen et al., 2021; Alibabaei et al., 2022). This finding also broadly agrees with Chapter 4 of this thesis, which found that fixed soil-moisture heuristics performed well compared to adaptive heuristics, which – given the simplicity and intuitive nature of these heuristics – should be seen as a good first step towards increasing productivity that producers would be more likely to adopt and implement.

Our finding that the value of DRL increases with the introduction of restrictive water use caps does provide interesting insights about where DRL is best implemented in real world cropping systems. As water resources become scarce and demands over them increase (Rosa et al., 2020; D'Odorico et al., 2020), water managers and policy makers may introduce restrictions on total water use – as some regions already do (Julio Berbel et al., 2019; Young et al., 2021). Our analysis indicates that DRL may perform well as water restrictions become more stringent and variable, with producers potentially being able to use these intelligent methods to mitigate impacts of increasing water insecurity. Our other finding that increased value in DRL for no-rain scenarios implies that DRL may also be better placed in arid areas where there is very little rainfall, and thus very little difference from year to year. By reducing the potential for unexpected weather scenarios, DRL agents will be free to overfit to the training data, as the validation data will be near identical.

Although we have been able to evaluate different action sets, decision frequencies, water use restrictions and climate variability, we have neglected comparing other important design choices such as different DRL algorithms, hyperparamaters, neural network architectures and input observations. For example, using larger neural networks or expanding the number of input observations will lead to more complex irrigation strategies, and hence potentially increased profits, though at the expense of longer training times and higher risk of overfitting. Investigating the effect of these and other DRL design choices in the value of DRL will be an important area of future work.

Applying DRL to irrigation scheduling, as we have done in this analysis, producers a number of issues for DRL that are not seen to the same extent in the typical environments DRL has been previously applied in. Assuming deployment within a real-world field, the agent will have to make decisions for several months before receiving a reward signal (e.g., profits calculated after harvesting). This issue is partly circumvented by training the agent within crop-simulation environments, which can simulate whole seasons within seconds. However, the costly and time consuming work of field trials will still have to be conducted eventually to test if model-optimised strategies will work as intended. This challenge of transferring simulated learning to real-world systems is receiving attention within robotics research (Zhao et al., 2020). However, the months needed to evaluate a single irrigation strategy is unique to irrigation scheduling. Another inherent difficulty of applying DRL to irrigation

scheduling is that there are an arbitrarily large number of environment configurations (e.g., crop, soil, climate, management constrains) that dramatically impact optimal decision-making. This issue either requires separate DRL agents to be trained on each configuration – which may be too computationally expensive or time consuming, as with optimised heuristics – or instead training one agent to make decisions for all configurations – which may be too difficult to learn good strategies. Moreover, for farmers to adopt a DRL irrigation strategy, the added value will have to be demonstrated for their particular configuration, the value of DRL may have be demonstrated in the field for many configurations in order for producers to adopt this new scheduling approach. Therefore this issue of configuration agnostic DRL irrigation strategies will be an essential area of future work.

Similarly, the irrigation scheduling problem specified in this work, and much of the previous literature (Cai et al., 2011; Linker et al., 2018; Jamal et al., 2019; Yang et al., 2020; Chen et al., 2021; Alibabaei et al., 2022), centres around scheduling irrigation for a single homogenous field over a number of seasons. This specification, although simple to implement, does favour an optimised heuristic approach (i.e., optimization algorithms can comfortably find a set of values that maximise average profits in one field over a set of *N* seasons). However, heuristics do not scale well with increasing complexity, for example, if the seasonal water limit changes during the season, the fixed heuristics would have to remain sub-optimal or be re-optimised each time the situation changes. Another example is in the case of multiple fields or crops, where decisions may have to be made about which field or crop to prioritise with the limited water supply available. Finally, an example from our case study region is dealing with a 5-year allocation of irrigation water supply, meaning decisions have to be made in the knowledge of how much water needs to be kept for future years. In these more complex – and more realistic – modelling scenarios, it is possible that DRL will demonstrate increased additional value.

To support research into these areas of future work, we have created *aquacrop-gym*, an open source python library that supports the training of DRL agents and comparisons between irrigation strategies within AquaCrop-OSPy. Training materials have been developed that guide the user to training their own DRL agents and duplicate the analysis presented in this work. The intention of this library is to be an actively developed framework that future researchers can test new irrigation strategies, and compare them to the current

best approaches. This framework will accelerate progress in optimised and intelligent irrigation scheduling research, as well as how rapidly this research can be converted to production.

5.5 Conclusions

Due to the increase in water scarcity issues globally, researchers are looking for innovative solutions to increasing agricultural water productivity. One of the methods being investigated is Deep Reinforcement Learning (DRL) which has had recent success inside complex decision making environments. In this article we assess the value of DRL for irrigation scheduling in comparison to optimised heuristics for a case study of irrigated maize production in southwest Nebraska. By designing a framework for DRL agents to interact with the AquaCrop-OSPy crop-simulation model, DRL agents were trained and evaluated over separate sets of climate data. Under normal climate variability and unrestricted irrigation, DRL failed to increase profits on the validation years compared to optimised soil-moisture thresholds. When rainfall was removed, and hence the biggest source of climate variability, DRL was able to surpass the performance soil-moisture thresholds. With the addition of restrictive irrigation caps, DRL was also able to achieve higher profits on the validation years than optimised thresholds. Our work firstly highlights that DRL is a promising avenue for improving irrigation scheduling, but DRL's additional value over optimised heuristics has yet to be demonstrated for all cases, especially given their increased complexity and data requirements. Conversely, optimised heuristics – given their relative simplicity – should be seen as an important first step to improving water productivity, before making large investments in AI based approaches.

Acknowledgments, Samples, and Data

The work contained in this article was funded by the National Environmental Research Council's Understanding the Earth, Atmosphere, and Ocean Doctoral Training Programme, Grant NE/L002469/1.

To support extensions of our research in line with future research directions outlined above, all code used in the analysis will be available in a public github repository (https://github.com/aquacropos/aquacrop-gym), allowing anyone to duplicate the analysis.

SUPPORTING INFORMATION

This supporting information includes Figure S1 and Table S1. Figure S1 compares historical and generated weather variables. Table S1 displays some of the important DRL hyperparameters.

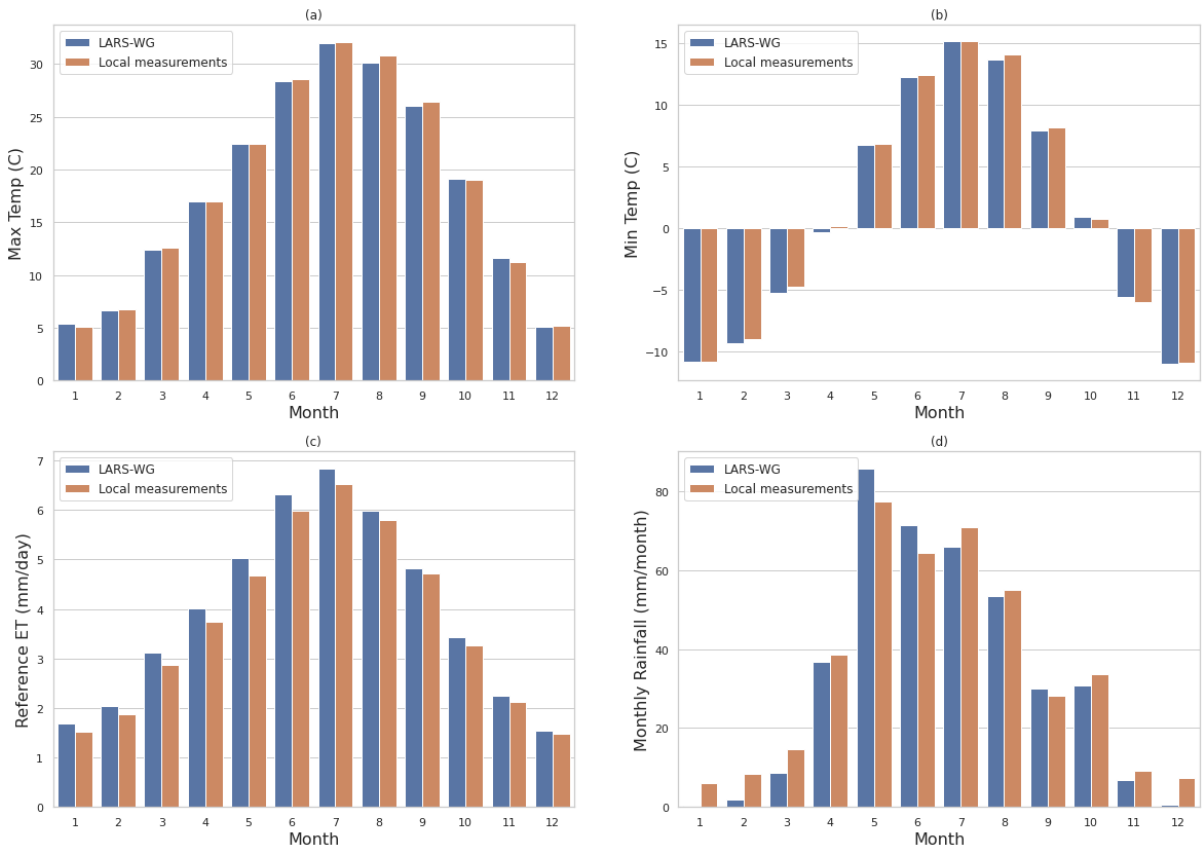

Figure 5.S1. The average daily (a) maximum temperature, (b) minimum temperature, (c) reference evapotranspiration, and (d) average monthly rainfall grouped by month. All figures compare the distribution of these weather variables between 100 years of LARS-WG generated data, and the local historical data that was originally imputed into LARS-WG.

Table 5.S1. Overview of DRL hyperparameters.

PARAMETER	VALUE
Network size	4 fully-connected layers (128 size)
Learning rate	5e-5
Training batch size	2056
Discount factor	1.
Number of CPU workers	8

6 CONCLUSIONS

As the largest sectoral user of freshwater globally, improving the productivity of agricultural water use (i.e., more crop per drop) is at the heart of efforts to address water scarcity, insecurity, and conflict. Improving irrigation scheduling practices is a key solution to improving water productivity, especially in currently high-productivity agricultural regions in developed countries where the potential gains from improving irrigation application technologies have largely been exhausted. Contributing toward this area of research, this thesis has assessed the value of improved information and management strategies for optimal irrigation scheduling. The first article presented AquaCrop-OSPy, which was developed to support the analysis within the thesis. The second article assessed the importance of reducing soil-moisture uncertainty in terms of reducing water use and increasing profits. The third article then assessed the value of adapting irrigation strategies in-season. Finally, the fourth article assessed the value of using Deep Reinforcement Learning for irrigation scheduling. This section will discuss the main findings and implications from these articles (Section 6.1), as well as areas for future work (Section 6.2).

6.1 KEY FINDINGS

One of the assumptions underpinning precision irrigation research is that providing farmers with more accurate information will lead to substantial increases in water productivity. However, an assessment of how information uncertainty impacts water use and farm profits – assuming realistic farmer behaviour – has not been previously performed. In response, Chapter 3 found that near-optimal irrigation scheduling decisions can be made without perfect information about soil moisture during the growing season. Results showed that it was the optimality of farmers' irrigation scheduling heuristics – rather than the accuracy of the information these heuristics are implemented based upon – that had the largest impact on water use and profits. These results are broadly consistent with previous research on the value of providing farmers with accurate soil and weather information (Bosch & Eidman, 1987; Botes et al., 1996; Fafchamps & Minten, 2012), as well as research by Linker & Kisekka (2017), who used a similar modelling approach to show that perfect real-time soil-moisture measurements may not be required to implement deficit irrigation strategies. The minimal increase in water use found in Chapter 3 is lower than that found in El Chami et al. (2019),

who found that realistic levels of soil-texture error lead to a 23% increase in water use for their case study of onion production in England. Differences such as their soil-moisture uncertainty purposefully being biased towards under-estimation of soil moisture – whereas the error in Chapter 3 was mean-zero – could partially explain those differences. Assuming the findings of Chapter 3 hold for other production environments, they suggest that efforts to improve irrigation water productivity should focus primarily on providing farmers with better advice on optimal irrigation scheduling, or designing tools to enable them to develop and test improved irrigation strategies for their specific cropping system.

Having to optimise irrigation strategies under future weather uncertainty will likely be a continual issue of producers globally, in particular given increasingly erratic rainfall in many regions caused by climate change (Turral et al., 2010; IPCC, 2022). Adaptive irrigation scheduling, whereby strategies are re-optimised within the season to account for the unfolding weather and future forecasts, has therefore been proposed as a key solution to improving the productivity and profitability of irrigation. However, it is unclear whether such adaptive irrigation strategies can significantly outperform simpler optimised fixed management heuristics (i.e., where scheduling rules do not change each year). Addressing this research gap, Chapter 4 found that in-season adaptation of irrigation scheduling rules was able to increase both average and minimum seasonal profits. However, this increase in average profit was marginal despite the additional computational and data requirements associated with implementing adaptive scheduling. The marginal effect of adaptation was due, in part, to the potential for maladaptive decisions to occur given uncertainty about future weather events. These results firstly agree with part of the current literature, finding irrigation strategies based on in-season stochastic re-optimization results in profits close to the maximum achievable (e.g., Wang & Cai, 2009; Cai et al., 2011; Hejazi et al., 2014; Jamal et al., 2018, 2019; Linker, 2021). The additional insights provided beyond these previous studies are that if irrigation strategies were not adjusted in-season, there would be only marginal decreases in profit. These findings imply that making near-optimal irrigation decisions under weather uncertainty can be achieved without the additional computational and data requirements of adaptive scheduling methods, assuming the fixed irrigation strategy is optimised for the expected variability in weather conditions.

Finally, the thesis explored the potential opportunities for emerging AI-driven approaches to improving irrigation scheduling practices beyond simpler optimised heuristics currently used by researchers and farmers. The potential advantage of AI approaches compared to simple rule-based heuristics is that AI scheduling can incorporate almost any available information (e.g., numerical, images, remote sensing) into complex strategies that can automatically adapt during the season. Chapter 5 aimed to compare one of these approaches (Deep Reinforcement Learning) to an optimised heuristic and found that Deep Reinforcement Learning performed best when irrigation was the dominant source of water for crop growth (i.e., when rainfall and its variability were limited), or when farmers faced restrictive limits on irrigation water use during the season. The results of this analysis firstly agree with previous work that DRL can learn complex irrigation strategies, and thus is a promising future method of irrigation scheduling (Yang et al., 2020; Chen et al., 2021; Alibabaei et al., 2022). However, by not evaluating DRL across the full climate variability, and against optimised heuristics, this previous literature may have overstated the value of DRL. The findings presented in Chapter 5 imply that for regions with significant year-to-year rainfall variability, new scheduling methods such as DRL may not provide much additional value compared with optimised heuristics. However, with increasing restrictions on irrigation water use (Young et al., 2021), farmers may turn to methods such as DRL to extract the maximum value from that limited water supply.

One common theme throughout the thesis has been the strong performance of scheduling irrigation using simple rule-based soil-moisture thresholds, which have been optimised to maximise profits over a range of historic weather years. This method of irrigation scheduling has been shown in the thesis to be robust to information uncertainty and generate comparable returns (in terms of productivity and profitability) compared with more complex adaptive and AI based irrigation scheduling. This conclusion is a significant contribution to the literature as it shows that large improvements in water productivity can be achieved without extensive investment in equipment, software and monitoring technologies required by advanced data-driven precision irrigation approaches. Specifically, the findings presented in this thesis suggest that improving agricultural water productivity is as much, if not more, dependent on incentivising change in how technologies are operationalised by farmers, rather than introducing new technologies or tools that have previously been the primary

focus of precision irrigation farming interventions (Adeyemi et al., 2017). As additional costs are one of the main barriers to adoption of water-saving technologies (USDA-NASS, 2018), improved irrigation management advice provides a low-cost pathway to improving water productivity and hence is likely to have greater traction with farmers. As previous research has found large heterogeneity in irrigation water use across farmers operating in similar production systems and using similar irrigation delivery systems (Foster et al., 2019; Gonçalves et al., 2020), improved scheduling advice may have the potential to close the gap between the less efficient and more efficient farmers.

6.2 FUTURE WORK

The findings presented within this thesis have highlighted a number of interesting and promising areas of future work. Firstly, field trials are required to evaluate model-optimised strategies (Section 6.2.1). Secondly, further research is required to gain a better understanding of how farmers currently make irrigation decisions (Section 6.2.2). Third, analysis needs to be extended beyond our case study of Nebraska maize, as well as to more complex modelling scenarios (Section 6.2.3). Finally, the tools developed within the thesis can be used to investigate more than field-level irrigation scheduling (Section 6.2.4).

6.2.1 Assessing performance in real-field conditions

Taghvaeian et al. (2020) highlighted the importance of *demonstrating the effectiveness of scientific irrigation scheduling* for researchers seeking to improve real-world agricultural water productivity. In this context, incentivising farmers to adopt the types of monitoring tools or optimised strategies used in this thesis will require farmer engagement to demonstrate these approaches are superior to the famers' current practises under real-field conditions, as opposed to in an idealised model environment. For example, the dual-simulation framework developed in Chapter 3 could be used to make irrigation scheduling decisions for real farmers' fields, with assimilation of in situ soil-moisture measurements to re-calibrate the model during the season to ensure the model accurately represents the current soil-moisture conditions (e.g., Linker & Ioslovich, 2017). Comparing this model-based scheduling to adjacent fields that have been irrigated according to the farmers' current practises would provide insights into the value of optimised scheduling methods in real production environments. Experiments such as this may find that optimised scheduling

leads to improvements in water productivity and profits, or conversely that farmers' current scheduling practises are already near-optimal, when technical, regulatory and environmental constraints and uncertainties they face are taken into account.

6.2.2 Improving understanding of farmer decision-making

Alongside real-world tests, there is also a need to improve scientific understanding of how farmers currently schedule irrigation and how scheduling choices are influenced by factors such as future weather forecasts (short term or seasonal), feedbacks from soil and crop monitoring, in-season changes in the water availability or irrigation costs, changes in crop market prices and information from neighbouring fields and farmers. Information on scheduling practices for the study areas in this thesis is largely derived from national surveys which only indicate the sources of information used by farmers (USDA-NASS, 2018), rather than how exactly this information influences decisions, as well as the relative importance placed on each of the different sources. This missing information could be gathered via surveys that are specifically targeted to asking how farmers schedule irrigation, or by collecting remote sensing data on variables already known to be important to decision making such as soil-moisture (Yinglan et al., 2022), or crop development (Tenreiro et al., 2021). By also estimating when irrigation is likely to have occurred (Le Page et al., 2020; Foster et al., 2020), decision tree algorithms can be used to predict the timing of irrigation events based on the remote sensing inputs (Perea et al., 2019). Potentially just as valuable as the predictive model, would be the knowledge gained from inspecting the resulting decision tree, giving insights into the relative importance of each of the input variables. Improved understanding of farmers' current irrigation strategies will enable researchers to more accurately represent current practises and hence more accurately evaluate the additional value of new scientific scheduling methods, such as optimised heuristics or DRL.

6.2.3 Testing performance in complex production environments

Analyses presented in this thesis focused primarily on case studies of field crop production in the central United States. However, the value of accurate monitoring or improved strategies may differ in another case study environments. For example, Chapter 5 found that the value of DRL increased when rainfall quantity and variability were reduced. This result motivates the expansion of the analysis presented in that chapter to other important agricultural case studies, to evaluate if that relationship holds, for example, in arid regions

with very little rainfall. Similarly, the optimised strategies developed in this thesis are optimal only for the weather variability seen during optimization or training. Climate shifts and changes necessitate that strategies be re-optimised periodically to maintain that future weather does not differ radically from expected. Therefore, future research should be explored that assesses how well optimised strategies perform under changing or nonstationary climate conditions, including how often strategies may need to be re-optimised between seasons as well as how many historic weather years should be used to best approximate the upcoming year.

Much of the work within the thesis has centred on irrigating a single crop on a single homogenous field. By increasing the complexity of the environment (e.g., by increasing the number of fields and/or crops), it is possible that optimised heuristics will no longer be able to achieve near-optimal performance, as adding further rules about which field or crop to prioritise may add too many decision variables to reliably identify global optimum heuristics through conventional rule-based optimization approaches alone. Conversely, DRL may scale well with the increasing complexity of the decision-making problem, as the number of input observations or the size of the internal neural network can simply scale to suit the complex environment. In more realistic and complex modelling scenarios such as this, DRL may offer greater value towards improving water productivity and profits, in comparison to the simpler environments used in this thesis.

6.2.4 From field to basin scale water productivity

Finally, the flexibility of tools developed in this thesis, in particular AquaCrop-OSPy, can contribute to several areas of future research that extend beyond the focus of this thesis (i.e., improving field-level irrigation scheduling). Field-level improvements in water productivity can sometimes have negative externalities when aggregated at basin scales. For example, water that is no longer lost through deep percolation or surface runoff may then no longer contribute to recharge to aquifers or return flows supporting downstream users or ecosystems (Van der Kooij et al., 2013; Berbel & Mateos, 2014; Berbel et al., 2019). Integration of AquaCrop-OSPy with basin-scale land-use and hydrologic models would allow more accurate estimates of the implications of changing irrigation management practices and technologies for regional water-use dynamics, for example, calculating how new water-use restrictions will effect streamflow depletion or groundwater return flows.

7 REFERENCES

Abdulkadri, A. O. (2003). Estimating risk aversion coefficients for dry land wheat, irrigated corn and dairy producers in Kansas. *Applied Economics*, *35*(7), 825–834. https://doi.org/10.1080/0003648032000050612

Abedinpour, M., Sarangi, A., Rajput, T. B. S., Singh, M., Pathak, H., & Ahmad, T. (2012). Performance evaluation of AquaCrop model for maize crop in a semi-arid environment. *Agricultural Water Management*, *110*, 55–66. https://doi.org/10.1016/j.agwat.2012.04.001

Abioye, E. A., Abidin, M. S. Z., Mahmud, M. S. A., Buyamin, S., Ishak, M. H. I., Rahman, M. K. I. A., Otuoze, A. O., Onotu, P., & Ramli, M. S. A. (2020). A review on monitoring and advanced control strategies for precision irrigation. *Computers and Electronics in Agriculture*, *173*, 105441. https://doi.org/10.1016/J.COMPAG.2020.105441

Adeyemi, O., Grove, I., Peets, S., Domun, Y., & Norton, T. (2018). Dynamic Neural Network Modelling of Soil Moisture Content for Predictive Irrigation Scheduling. *Sensors*, *18*(10), 3408. https://doi.org/10.3390/s18103408

Adeyemi, O., Grove, I., Peets, S., & Norton, T. (2017). Advanced Monitoring and Management Systems for Improving Sustainability in Precision Irrigation. *Sustainability*, *9*(3), 353. https://doi.org/10.3390/su9030353

Agarwal, A., Babel, M. S., & Maskey, S. (2014). Analysis of future precipitation in the Koshi river basin, Nepal. *Journal of Hydrology*, *513*, 422–434. https://doi.org/10.1016/J.JHYDROL.2014.03.047

Al-Jamal, M. S., Ball, S., & Sammis, T. W. (2001). Comparison of sprinkler, trickle and furrow irrigation efficiencies for onion production. *Agricultural Water Management*, *46*(3), 253–266. https://doi.org/10.1016/S0378-3774(00)00089-5

Alibabaei, K., Gaspar, P. D., Assunção, E., Alirezazadeh, S., & Lima, T. M. (2022). Irrigation optimization with a deep reinforcement learning model: Case study on a site in Portugal. *Agricultural Water Management*, *263*, 107480. https://doi.org/10.1016/j.agwat.2022.107480

Allen, R. G., Pereira, L. S., Raes, D., & Smith, M. (1998). Crop evapotranspiration-Guidelines for computing crop water requirements-FAO Irrigation and drainage paper 56. In *Fao* (Vol. 9).

Andrychowicz, O. A. M., Baker, B., Chociej, M., Józefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., Schneider, J., Sidor, S., Tobin, J., Welinder, P., Weng, L., & Zaremba, W. (2019). Learning dexterous in-hand manipulation: *Https://Doi.Org/10.1177/0278364919887447*, *39*(1), 3–20. https://doi.org/10.1177/0278364919887447

Araya, A., Kisekka, I., & Holman, J. (2016). Evaluating deficit irrigation management strategies for grain sorghum using AquaCrop. *Irrigation Science*, *34*(6), 465–481. https://doi.org/10.1007/s00271-016-0515-7

Araya, A., Kisekka, I., Vara Prasad, P. V., & Gowda, P. H. (2017). Evaluating Optimum Limited Irrigation Management Strategies for Corn Production in the Ogallala Aquifer Region. *Journal of Irrigation and Drainage Engineering*, *143*(10), 04017041. https://doi.org/10.1061/(asce)ir.1943-4774.0001228

Attia, A., Rajan, N., Xue, Q., Nair, S., Ibrahim, A., & Hays, D. (2016). Application of DSSAT-CERES-Wheat model to simulate winter wheat response to irrigation management in the Texas High Plains. *Agricultural Water Management*, *165*, 50–60.

Awal, R., Bayabil, H. K., & Fares, A. (2016). Analysis of Potential Future Climate and Climate Extremes in the Brazos Headwaters Basin, Texas. *Water 2016, Vol. 8, Page 603*, *8*(12), 603. https://doi.org/10.3390/W8120603

Badia, A. P., Piot, B., Kapturowski, S., Sprechmann, P., Vitvitskyi, A., Guo, D., & Blundell, C. (2020). Agent57: Outperforming the Atari Human Benchmark. *37th International Conference on Machine Learning, ICML 2020*, *PartF168147-1*, 484–494. https://doi.org/10.48550/arxiv.2003.13350

Bartlett, A. C., Andales, A. A., Arabi, M., & Bauder, T. A. (2015). A smartphone app to extend use of a cloud-based irrigation scheduling tool. *Computers and Electronics in Agriculture*, *111*, 127–130. https://doi.org/10.1016/j.compag.2014.12.021

Berbel, J., & Mateos, L. (2014). Does investment in irrigation technology necessarily

generate rebound effects? A simulation analysis based on an agro-economic model. *Agricultural Systems*, *128*, 25–34. https://doi.org/10.1016/j.agsy.2014.04.002

Berbel, Julio, Expósito, A., Gutiérrez-Martín, C., & Mateos, L. (2019). Effects of the Irrigation Modernization in Spain 2002–2015. *Water Resources Management*, *33*(5), 1835–1849. https://doi.org/10.1007/s11269-019-02215-w

Blonquist, J. M., Jones, S. B., & Robinson, D. A. (2006). Precise irrigation scheduling for turfgrass using a subsurface electromagnetic soil moisture sensor. *Agricultural Water Management*, *84*(1–2), 153–165. https://doi.org/10.1016/j.agwat.2006.01.014

Bosch, D. J., & Eidman, V. R. (1987). Valuing Information When Risk Preferences Are Nonneutral: An Application to Irrigation Scheduling. *American Journal of Agricultural Economics*, *69*(3), 658–668. https://doi.org/10.2307/1241700

Bossio, D., Geheb, K., & Critchley, W. (2010). Managing water by managing land: Addressing land degradation to improve water productivity and rural livelihoods. *Agricultural Water Management*, *97*(4), 536–542. https://doi.org/10.1016/J.AGWAT.2008.12.001

Botes, J. H. F., Bosch, D. J., & Oosthuizen, L. K. (1996). A simulation and optimization approach for evaluating irrigation information. *Agricultural Systems*, *51*(2), 165–183. https://doi.org/10.1016/0308-521X(95)00042-4

Brocca, L., Morbidelli, R., Melone, F., & Moramarco, T. (2007). Soil moisture spatial variability in experimental areas of central Italy. *Journal of Hydrology*, *333*(2–4), 356–373. https://doi.org/10.1016/j.jhydrol.2006.09.004

Butler, J. J., Bohling, G. C., Whittemore, D. O., & Wilson, B. B. (2020). Charting Pathways Toward Sustainability for Aquifers Supporting Irrigated Agriculture. *Water Resources Research*, *56*(10), e2020WR027961. https://doi.org/10.1029/2020WR027961

Cai, X., Hejazi, M. I., & Wang, D. (2011). Value of Probabilistic Weather Forecasts: Assessment by Real-Time Optimization of Irrigation Scheduling. *Journal of Water Resources Planning and Management*, *137*(5), 391–403. https://doi.org/10.1061/(asce)wr.1943-5452.0000126

Çakir, R. (2004). Effect of water stress at different development stages on vegetative and

reproductive growth of corn. *Field Crops Research*, *89*(1), 1–16. https://doi.org/10.1016/j.fcr.2004.01.005

Camp, C. R. (1998). Subsurface drip irrigation: A review. In *Transactions of the American Society of Agricultural Engineers* (Vol. 41, Issue 5, pp. 1353–1367). American Society of Agricultural and Biological Engineers. https://doi.org/10.13031/2013.17309

Chaubey, I., Haan, C. T., Grunwald, S., & Salisbury, J. M. (1999). Uncertainty in the model parameters due to spatial variability of rainfall. *Journal of Hydrology*, *220*(1–2), 48–61. https://doi.org/10.1016/S0022-1694(99)00063-3

Chen, M., Cui, Y., Wang, X., Xie, H., Liu, F., Luo, T., Zheng, S., & Luo, Y. (2021). A reinforcement learning approach to irrigation decision-making for rice using weather forecasts. *Agricultural Water Management*, *250*, 106838. https://doi.org/10.1016/j.agwat.2021.106838

Closas, A., & Molle, F. (2018). Chronicle of a demise foretold: State vs. local groundwater management in Texas and the High Plains aquifer system. *Water Alternatives*, *11*(3), 511–532.

CropWatch. (2018). *2018 Corn Yield Forecast Table 2 Soil Types*. University of Nebraska–Lincoln. https://cropwatch.unl.edu/2018/2018-corn-yield-forecast-table-2-soil-types

CropWatch. (2019a). *Nebraska Crop Budgets*. https://cropwatch.unl.edu/budgets

CropWatch. (2019b). *Soil Moisture Sensor Programs Offered by Nebraska Natural Resources Districts*. University of Nebraska–Lincoln. https://cropwatch.unl.edu/soil-moisture-sensor-programs-offered-nebraska-natural-resources-districts

D'Odorico, P., Chiarelli, D. D., Rosa, L., Bini, A., Zilberman, D., & Rulli, M. C. (2020). The global value of water in agriculture. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(36), 21985–21993. https://doi.org/10.1073/PNAS.2005835117/SUPPL_FILE/PNAS.2005835117.SAPP.PDF

Daccache, A., Knox, J. W., Weatherhead, E. K., Daneshkhah, A., & Hess, T. M. (2015). Implementing precision irrigation in a humid climate – Recent experiences and on-going challenges. *Agricultural Water Management*, *147*, 135–143.

https://doi.org/10.1016/j.agwat.2014.05.018

Dale, A., Fant, C., Strzepek, K., Lickley, M., & Solomon, S. (2017). Climate model uncertainty in impact assessments for agriculture: A multi-ensemble case study on maize in sub-Saharan Africa. *Earth's Future*, *5*(3), 337–353. https://doi.org/10.1002/2017EF000539

Davis, S. L., & Dukes, M. D. (2010). Irrigation scheduling performance by evapotranspiration-based controllers. *Agricultural Water Management*, *98*(1), 19–28. https://doi.org/10.1016/J.AGWAT.2010.07.006

De Paly, M., & Zell, A. (2009). Optimal irrigation scheduling with evolutionary algorithms. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *5484 LNCS*, 142–151. https://doi.org/10.1007/978-3-642-01129-0_18

Deines, J. M., Schipanski, M. E., Golden, B., Zipper, S. C., Nozari, S., Rottler, C., Guerrero, B., & Sharda, V. (2020). Transitions from irrigated to dryland agriculture in the Ogallala Aquifer: Land use suitability and regional economic impacts. *Agricultural Water Management*, *233*, 106061. https://doi.org/10.1016/j.agwat.2020.106061

Dobriyal, P., Qureshi, A., Badola, R., & Hussain, S. A. (2012). A review of the methods available for estimating soil moisture and its implications for water resource management. *Journal of Hydrology*, *458–459*, 110–117. https://doi.org/10.1016/j.jhydrol.2012.06.021

Dong, J., Gruda, N., Lam, S. K., Li, X., & Duan, Z. (2018). Effects of elevated CO2 on nutritional quality of vegetables: A review. *Frontiers in Plant Science*, *9*, 924. https://doi.org/10.3389/FPLS.2018.00924/BIBTEX

Dukes, M. D., Zotarelli, L., & Morgan, K. T. (2010). Use of Irrigation Technologies for Vegetable Crops in Florida. *HortTechnology* , *20*(1), 133–142. http://horttech.ashspublications.org/content/20/1/133.abstract

Eberhart, R., & Kennedy, J. (1995). Particle swarm optimization. *Proceedings of the IEEE International Conference on Neural Networks*, *4*, 1942–1948.

El Chami, D., Knox, J. W., Daccache, A., & Weatherhead, E. K. (2019). Assessing the financial

and environmental impacts of precision irrigation in a humid climate. *Horticultural Science*, *46*(1), 43–52. https://doi.org/10.17221/116/2017-HORTSCI

Evans, R. G., & Sadler, E. J. (2008). Methods and technologies to improve efficiency of water use. *Water Resources Research*, *44*(7). https://doi.org/10.1029/2007WR006200

Expósito, A., & Berbel, J. (2019). Drivers of Irrigation Water Productivity and Basin Closure Process: Analysis of the Guadalquivir River Basin (Spain). *Water Resources Management*, *33*(4), 1439–1450. https://doi.org/10.1007/S11269-018-2170-7/TABLES/5

Fafchamps, M., & Minten, B. (2012). Impact of SMS-based agricultural information on Indian farmers. *World Bank Economic Review*, *26*(3), 383–414. https://doi.org/10.1093/wber/lhr056

Feki, M., Ravazzani, G., Ceppi, A., & Mancini, M. (2018). Influence of soil hydraulic variability on soil moisture simulations and irrigation scheduling in a maize field. *Agricultural Water Management*, *202*(August 2017), 183–194. https://doi.org/10.1016/j.agwat.2018.02.024

Fereres, E., & Soriano, M. A. (2006). Deficit irrigation for reducing agricultural water use. *Journal of Experimental Botany*, *58*(2), 147–159. https://doi.org/10.1093/jxb/erl165

Fernández García, I., Rodríguez Díaz, J. A., Camacho Poyato, E., Montesinos, P., & Berbel, J. (2014). Effects of modernization and medium term perspectives on water and energy use in irrigation districts. *Agricultural Systems*, *131*, 56–63. https://doi.org/10.1016/J.AGSY.2014.08.002

Foley, J. A., Ramankutty, N., Brauman, K. A., Cassidy, E. S., Gerber, J. S., Johnston, M., Mueller, N. D., O'Connell, C., Ray, D. K., West, P. C., Balzer, C., Bennett, E. M., Carpenter, S. R., Hill, J., Monfreda, C., Polasky, S., Rockström, J., Sheehan, J., Siebert, S., … Zaks, D. P. M. (2011). Solutions for a cultivated planet. *Nature 2011 478:7369*, *478*(7369), 337–342. https://doi.org/10.1038/nature10452

Foster, T., & Brozović, N. (2018). Simulating Crop-Water Production Functions Using Crop Growth Models to Support Water Policy Assessments. *Ecological Economics*, *152*(October 2017), 9–21. https://doi.org/10.1016/j.ecolecon.2018.05.019

Foster, T., Brozović, N., & Butler, A. P. (2014). Modeling irrigation behavior in groundwater systems. *Water Resources Research*, *50*(8), 6370–6389. https://doi.org/10.1002/2014WR015620

Foster, T., Brozović, N., & Butler, A. P. (2015). Analysis of the impacts of well yield and groundwater depth on irrigated agriculture. *Journal of Hydrology*, *523*, 86–96. https://doi.org/10.1016/j.jhydrol.2015.01.032

Foster, T., Brozović, N., & Butler, A. P. (2017). Effects of initial aquifer conditions on economic benefits from groundwater conservation. *Water Resources Research*, *53*(1), 744–762. https://doi.org/10.1002/2016WR019365

Foster, T., Brozović, N., Butler, A. P., Neale, C. M. U., Raes, D., Steduto, P., Fereres, E., & Hsiao, T. C. (2017). AquaCrop-OS: An open source version of FAO's crop water productivity model. *Agricultural Water Management*, *181*, 18–22. https://doi.org/10.1016/j.agwat.2016.11.015

Foster, T., Gonçalves, I. Z., Campos, I., Neale, C. M. U. U., & Brozović, N. (2019). Assessing landscape scale heterogeneity in irrigation water use with remote sensing and in situ monitoring. *Environmental Research Letters*, *14*(2), 024004. https://doi.org/10.1088/1748-9326/aaf2be

Foster, T., Mieno, T., & Brozović, N. (2020). Satellite-Based Monitoring of Irrigation Water Use: Assessing Measurement Errors and Their Implications for Agricultural Water Management Policy. *Water Resources Research*, *56*(11). https://doi.org/10.1029/2020WR028378

Freitas, D., Lopes, L. G., & Morgado-Dias, F. (2020). Particle Swarm Optimisation: A Historical Review Up to the Current Developments. *Entropy*, *22*(3), 362. https://doi.org/10.3390/e22030362

García-Vila, M., Fereres, E., Mateos, L., Orgaz, F., & Steduto, P. (2009). Deficit Irrigation Optimization of Cotton with AquaCrop. *Agronomy Journal*, *101*(3), 477–487. https://doi.org/10.2134/AGRONJ2008.0179S

García-Vila, Margarita, & Fereres, E. (2012). Combining the simulation crop model AquaCrop with an economic model for the optimization of irrigation management at farm level.

European Journal of Agronomy, *36*(1), 21–31. https://doi.org/10.1016/j.eja.2011.08.003

Geerts, S., Raes, D., & Garcia, M. (2010). Using AquaCrop to derive deficit irrigation schedules. *Agricultural Water Management*, *98*(1), 213–216. https://doi.org/10.1016/j.agwat.2010.07.003

Gonçalves, I. Z., Mekonnen, M. M., Neale, C. M. U., Campos, I., & Neale, M. R. (2020). Temporal and spatial variations of irrigation water use for commercial corn fields in Central Nebraska. *Agricultural Water Management*, *228*, 105924. https://doi.org/10.1016/j.agwat.2019.105924

González Perea, R., Camacho Poyato, E., Montesinos, P., & Rodríguez Díaz, J. A. (2019). Prediction of irrigation event occurrence at farm level using optimal decision trees. *Computers and Electronics in Agriculture*, *157*(December 2018), 173–180. https://doi.org/10.1016/j.compag.2018.12.043

Goosheh, M., Pazira, E., Gholami, A., Andarzian, B., & Panahpour, E. (2018). Improving Irrigation Scheduling of Wheat to Increase Water Productivity in Shallow Groundwater Conditions Using Aquacrop. *Irrigation and Drainage*, *67*(5), 738–754. https://doi.org/10.1002/ird.2288

Grafton, R. Q., Williams, J., Perry, C. J., Molle, F., Ringler, C., Steduto, P., Udall, B., Wheeler, S. A., Wang, Y., Garrick, D., & Allen, R. G. (2018). The paradox of irrigation efficiency. *Science*, *361*(6404), 748–750. https://doi.org/10.1126/SCIENCE.AAT9314/SUPPL_FILE/AAT9314_GRAFTON_SM.PDF

Gutierrez, J., Villa-Medina, J. F., Nieto-Garibay, A., & Porta-Gandara, M. A. (2014). Automated Irrigation System Using a Wireless Sensor Network and GPRS Module. *IEEE Transactions on Instrumentation and Measurement*, *63*(1), 166–176. https://doi.org/10.1109/TIM.2013.2276487

Haghverdi, A., Leib, B. G., Washington-Allen, R. A., Ayers, P. D., & Buschermohle, M. J. (2015). Perspectives on delineating management zones for variable rate irrigation. *Computers and Electronics in Agriculture*, *117*, 154–167. https://doi.org/10.1016/J.COMPAG.2015.06.019

Hazeleger, W., Wang, X., Severijns, C., Ştefănescu, S., Bintanja, R., Sterl, A., Wyser, K., Semmler, T., Yang, S., van den Hurk, B., van Noije, T., van der Linden, E., & van der Wiel, K. (2012). EC-Earth V2.2: Description and validation of a new seamless earth system prediction model. *Climate Dynamics*, *39*(11), 2611–2629. https://doi.org/10.1007/s00382-011-1228-5

Hedley, C. B., Roudier, P., Yule, I. J., Ekanayake, J., & Bradbury, S. (2013). Soil water status and water table depth modelling using electromagnetic surveys for precision irrigation scheduling. *Geoderma*, *199*, 22–29. https://doi.org/10.1016/J.GEODERMA.2012.07.018

Hejazi, M. I., Cai, X., Yuan, X., Liang, X.-Z., & Kumar, P. (2014). Incorporating Reanalysis-Based Short-Term Forecasts from a Regional Climate Model in an Irrigation Scheduling Optimization Problem. *Journal of Water Resources Planning and Management*, *140*(5), 699–713. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000365

Heng, L. K., Hsiao, T., Evett, S., Howell, T., & Steduto, P. (2009). Validating the FAO AquaCrop Model for Irrigated and Water Deficient Field Maize. *Agronomy Journal*, *101*(3), 488–498. https://doi.org/10.2134/agronj2008.0029xs

Hoerling, M., Eischeid, J., Kumar, A., Leung, R., Mariotti, A., Mo, K., Schubert, S., & Seager, R. (2014). Causes and predictability of the 2012 great plains drought. *Bulletin of the American Meteorological Society*, *95*(2), 269–282. https://doi.org/10.1175/BAMS-D-13-00055.1

Howell, T. A. (2001). Enhancing Water Use Efficiency in Irrigated Agriculture. *Agronomy Journal*, *93*(2), 281–289. https://doi.org/10.2134/AGRONJ2001.932281X

HPRCC. (2016). *High Plains Regional Climate Center.* https://hprcc.unl.edu/index.php

Hsiao, T. C., Heng, L., Steduto, P., Rojas-Lara, B., Raes, D., & Fereres, E. (2009). AquaCrop-The FAO Crop Model to Simulate Yield Response to Water: III. Parameterization and Testing for Maize. *Agronomy Journal*, *101*(3), 448–459. https://doi.org/10.2134/agronj2008.0218s

Ifft, J., Bigelow, D. P., & Savage, J. (2018). The impact of irrigation restrictions on cropland values in Nebraska. *Journal of Agricultural and Resource Economics*, *43*(2), 195–214. https://doi.org/10.22004/ag.econ.273446

Igbadun, H. E., Tarimo, A. K. P. R., Salim, B. A., & Mahoo, H. F. (2007). Evaluation of selected crop water production functions for an irrigated maize crop. *Agricultural Water Management*, *94*(1–3), 1–10. https://doi.org/10.1016/j.agwat.2007.07.006

Ihuoma, S. O., & Madramootoo, C. A. (2017). Recent advances in crop water stress detection. *Computers and Electronics in Agriculture*, *141*, 267–275. https://doi.org/10.1016/j.compag.2017.07.026

IPCC. (2022). *Climate Change 2022: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [H.-O. Pörtner, D.C. Roberts, M. Tignor, E.S. Poloczanska, K. Mintenbeck, A. Aleg*.

Jamal, A., Linker, R., & Housh, M. (2018). Comparison of Various Stochastic Approaches for Irrigation Scheduling Using Seasonal Climate Forecasts. *Journal of Water Resources Planning and Management*, *144*(7), 04018028. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000951

Jamal, A., Linker, R., & Housh, M. (2019). Optimal Irrigation with Perfect Weekly Forecasts versus Imperfect Seasonal Forecasts. *Journal of Water Resources Planning and Management*, *145*(5), 06019003. https://doi.org/10.1061/(ASCE)WR.1943-5452.0001066

Jiang, Y., Zhang, L., Zhang, B., He, C., Jin, X., & Bai, X. (2016). Modeling irrigation management for water conservation by DSSAT-maize model in arid northwestern China. *Agricultural Water Management*, *177*, 37–45. https://doi.org/10.1016/j.agwat.2016.06.014

Jones, H. G. (2004). Irrigation scheduling: advantages and pitfalls of plant-based methods. *Journal of Experimental Botany*, *55*(407), 2427–2436. https://doi.org/10.1093/jxb/erh213

Jones, J. ., Hoogenboom, G., Porter, C. ., Boote, K. ., Batchelor, W. ., Hunt, L. ., Wilkens, P. ., Singh, U., Gijsman, A. ., & Ritchie, J. . (2003). The DSSAT cropping system model. *European Journal of Agronomy*, *18*(3–4), 235–265. https://doi.org/10.1016/S1161-0301(02)00107-7

Jones, J. W., Hansen, J. W., Royce, F. S., & Messina, C. D. (2000). Potential benefits of climate forecasting to agriculture. *Agriculture, Ecosystems & Environment*, *82*(1–3), 169–184. https://doi.org/10.1016/S0167-8809(00)00225-5

Jozsef Szilagyi. (2000). Streamflow Depletion Investigations in the Republican River Basin: Colorado, Nebraska, and Kansas. *Journal of Environmental Systems*, *27*(3), 251–263. https://doi.org/10.2190/0MT9-V901-N4UQ-N7UM

Kahil, M. T., Connor, J. D., & Albiac, J. (2015). Efficient water management policies for irrigation adaptation to climate change in Southern Europe. *Ecological Economics*, *120*, 226–233. https://doi.org/10.1016/j.ecolecon.2015.11.004

Kang, S., Shi, W., & Zhang, J. (2000). An improved water-use efficiency for maize grown under regulated deficit irrigation. *Field Crops Research*, *67*(3), 207–214. https://doi.org/10.1016/S0378-4290(00)00095-2

Kelly, T. D., & Foster, T. (2021). AquaCrop-OSPy: Bridging the gap between research and practice in crop-water modeling. *Agricultural Water Management*, *254*, 106976. https://doi.org/10.1016/j.agwat.2021.106976

Kelly, T. D., Foster, T., Schultz, D. M., & Mieno, T. (2021). The effect of soil-moisture uncertainty on irrigation water use and farm profits. *Advances in Water Resources*, *154*, 103982. https://doi.org/10.1016/j.advwatres.2021.103982

Kloss, S., Grundmann, J., Seidel, S. J., Werisch, S., Trümmner, J., Schmidhalter, U., & Schütze, N. (2014). Investigation of deficit irrigation strategies combining SVAT-modeling, optimization and experiments. *Environmental Earth Sciences*, *72*(12), 4901–4915. https://doi.org/10.1007/s12665-014-3463-7

Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., Abdalla, S., & Willing, C. (2016). Jupyter Notebooks—a publishing format for reproducible computational workflows. *Positioning and Power in Academic Publishing: Players, Agents and Agendas - Proceedings of the 20th International Conference on Electronic Publishing, ELPUB 2016*, 87–90. https://doi.org/10.3233/978-1-61499-649-1-87

Koundouri, P., Nauges, C., & Tzouvelekas, V. (2006). Technology Adoption under Production

Uncertainty: Theory and Application to Irrigation Technology. *American Journal of Agricultural Economics*, *88*(3), 657–670. https://doi.org/10.1111/J.1467-8276.2006.00886.X

Kukal, M. S., Irmak, S., & Sharma, K. (2019). Development and Application of a Performance and Operational Feasibility Guide to Facilitate Adoption of Soil Moisture Sensors. *Sustainability*, *12*(1), 321. https://doi.org/10.3390/su12010321

Le Page, M., Jarlan, L., El Hajj, M. M., Zribi, M., Baghdadi, N., & Boone, A. (2020). Potential for the Detection of Irrigation Events on Maize Plots Using Sentinel-1 Soil Moisture Products. *Remote Sensing 2020, Vol. 12, Page 1621*, *12*(10), 1621. https://doi.org/10.3390/RS12101621

Leathes, W., Knox, J. W., Kay, M. G., Trawick, P., & Rodriguez-Diaz, J. A. (2008). Developing UK farmers' institutional capacity to defend their water rights and effectively manage limited water resources. *Irrigation and Drainage*, *57*(3), 322–331. https://doi.org/10.1002/IRD.436

Lehmann, N., Briner, S., & Finger, R. (2013). The impact of climate and price risks on agricultural land use and crop management decisions. *Land Use Policy*, *35*, 119–130. https://doi.org/10.1016/j.landusepol.2013.05.008

Li, F., Yu, D., & Zhao, Y. (2019). Irrigation Scheduling Optimization for Cotton Based on the AquaCrop Model. *Water Resources Management*, *33*(1), 39–55. https://doi.org/10.1007/s11269-018-2087-1

Li, Y., Guan, K., Schnitkey, G. D., DeLucia, E., & Peng, B. (2019). Excessive rainfall leads to maize yield loss of a comparable magnitude to extreme drought in the United States. *Global Change Biology*, *25*(7), gcb.14628. https://doi.org/10.1111/gcb.14628

Liang, E., Liaw, R., Moritz, P., Nishihara, R., Fox, R., Goldberg, K., Gonzalez, J. E., Jordan, M. I., & Stoica, I. (2017). RLlib: Abstractions for Distributed Reinforcement Learning. *35th International Conference on Machine Learning, ICML 2018*, *7*, 4768–4780. https://doi.org/10.48550/arxiv.1712.09381

Linker, R., & Sylaios, G. (2016). Efficient model-based sub-optimal irrigation scheduling using imperfect weather forecasts. *Computers and Electronics in Agriculture*, *130*, 118–127.

https://doi.org/10.1016/j.compag.2016.10.004

Linker, R., Sylaios, G., Tsakmakis, I., Ramos, T., Simionesei, L., Plauborg, F., & Battilani, A. (2018). Sub-optimal model-based deficit irrigation scheduling with realistic weather forecasts. *Irrigation Science*, *36*(6), 349–362. https://doi.org/10.1007/s00271-018-0592-x

Linker, Raphael. (2021). Stochastic model-based optimization of irrigation scheduling. *Agricultural Water Management*, *243*, 106480. https://doi.org/10.1016/J.AGWAT.2020.106480

Linker, Raphael, & Ioslovich, I. (2017). Assimilation of canopy cover and biomass measurements in the crop model AquaCrop. *Biosystems Engineering*, *162*, 57–66. https://doi.org/10.1016/J.BIOSYSTEMSENG.2017.08.003

Linker, Raphael, Ioslovich, I., Sylaios, G., Plauborg, F., & Battilani, A. (2016). Optimal model-based deficit irrigation scheduling using AquaCrop: A simulation study with cotton, potato and tomato. *Agricultural Water Management*, *163*, 236–243. https://doi.org/10.1016/j.agwat.2015.09.011

Linker, Raphael, & Kisekka, I. (2017). Model-Based Deficit Irrigation of Maize in Kansas. *Transactions of the ASABE*, *60*(6), 2011–2022. https://doi.org/10.13031/trans.12341

Loch, A., Pérez-Blanco, C. D., Carmody, E., Felbab-Brown, V., Adamson, D., & Seidl, C. (2020). Grand theft water and the calculus of compliance. *Nature Sustainability*, *3*(12), 1012–1018. https://doi.org/10.1038/s41893-020-0589-3

López-Morales, C., & Duchin, F. (2011). POLICIES AND TECHNOLOGIES FOR A SUSTAINABLE USE OF WATER IN MEXICO: A SCENARIO ANALYSIS. *Economic Systems Research*, *23*(4), 387–407. https://doi.org/10.1080/09535314.2011.635138

Maisiri, N., Senzanje, A., Rockstrom, J., & Twomlow, S. J. (2005). On farm evaluation of the effect of low cost drip irrigation on water and crop productivity compared to conventional surface irrigation system. *Physics and Chemistry of the Earth, Parts A/B/C*, *30*(11–16), 783–791. https://doi.org/10.1016/J.PCE.2005.08.021

Martínez-Romero, A., Domínguez, A., & Landeras, G. (2019). Regulated deficit irrigation

strategies for different potato cultivars under continental Mediterranean-Atlantic conditions. *Agricultural Water Management*, *216*, 164–176. https://doi.org/10.1016/j.agwat.2019.01.030

Mathworks. (n.d.). *Particle Swarm Optimization Algorithm*. Retrieved May 11, 2020, from https://uk.mathworks.com/help/gads/particle-swarm-optimization-algorithm.html#mw_522b9230-864b-47d1-a0db-1bf6c882d862

Mavrovouniotis, M., Li, C., & Yang, S. (2017). A survey of swarm intelligence for dynamic optimization: Algorithms and applications. *Swarm and Evolutionary Computation*, *33*, 1–17. https://doi.org/10.1016/j.swevo.2016.12.005

McCarthy, A. C., Hancock, N. H., & Raine, S. R. (2010). VARIwise: A general-purpose adaptive control simulation framework for spatially and temporally varied irrigation at sub-field scale. *Computers and Electronics in Agriculture*, *70*(1), 117–128. https://doi.org/10.1016/j.compag.2009.09.011

McGuire, V.L. (2017). Water-Level and Recoverable Water in Storage Changes, High Plains Aquifer, Predevelopment to 2015 and 2013-15. *U.S. Geological Survey Scientific Investigations Report 2017–5040*, 14. https://doi.org/10.3133/sir20175040

McGuire, Virginia L. (2014). Water-level changes and change in water in storage in the High Plains aquifer, predevelopment to 2013 and 2011-13. In *Scientific Investigations Report*. https://doi.org/10.3133/SIR20145218

Mekonnen, M. M., & Hoekstra, A. Y. (2016). Four billion people facing severe water scarcity. *Science Advances*, *2*(2), e1500323. https://doi.org/10.1126/sciadv.1500323

Menapace, L., Colson, G., & Raffaelli, R. (2013). Risk Aversion, Subjective Beliefs, and Farmer Risk Management Strategies. *American Journal of Agricultural Economics*, *95*(2), 384–389. https://doi.org/10.1093/ajae/aas107

Molden, D., & FAO. (2020). The State of Food and Agriculture 2020. In D. Molden (Ed.), *International Water Management Institute*. FAO. https://doi.org/10.4060/cb1447en

Montginoul, M., Rinaudo, J. D., Brozović, N., & Donoso, G. (2016). Controlling groundwater exploitation through economic instruments: Current practices, challenges and

innovative approaches. In *Integrated Groundwater Management: Concepts, Approaches and Challenges* (pp. 551–581). Springer International Publishing. https://doi.org/10.1007/978-3-319-23576-9_22

Morison, J. I. ., Baker, N. ., Mullineaux, P. ., & Davies, W. . (2008). Improving water use in crop production. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1491), 639–658. https://doi.org/10.1098/rstb.2007.2175

Mwangi, M., & Kariuki, S. (2015). Factors Determining Adoption of New Agricultural Technology by Smallholder Farmers in Developing Countries. *Journal of Economic and Sustainable Development*, *6*(5), 208–216. www.iiste.org

Najafi, P., & Tabatabaei, S. H. (2007). Effect of using subsurface drip irrigation and ET-HS model to increase WUE in irrigation of some crops. *Irrigation and Drainage*, *56*(4), 477–486. https://doi.org/10.1002/IRD.322

Nielsen, D. R., Biggar, J. W., & Erh, K. T. (1973). Spatial variability of field-measured soil-water properties. *Hilgardia*, *42*(7), 215–259. https://doi.org/10.3733/hilg.v42n07p215

Noory, H., Liaghat, A. M., Parsinejad, M., & Haddad, O. B. (2012). Optimizing Irrigation Water Allocation and Multicrop Planning Using Discrete PSO Algorithm. *Journal of Irrigation and Drainage Engineering*, *138*(5), 437–444. https://doi.org/10.1061/(ASCE)IR.1943-4774.0000426

Nuccio, M. L., Paul, M., Bate, N. J., Cohn, J., & Cutler, S. R. (2018). Where are the drought tolerant crops? An assessment of more than two decades of plant biotechnology effort in crop improvement. *Plant Science*, *273*, 110–119. https://doi.org/10.1016/J.PLANTSCI.2018.01.020

OBrien, D. M., Rogers, D. H., Lamm, F. R., & Clark, G. A. (1998). An economic comparison of subsurface drip and center pivot sprinkler irrigation systems. *Applied Engineering in Agriculture*, *14*(4), 391–398.

OpenAI, :, Berner, C., Brockman, G., Chan, B., Cheung, V., Dębiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., Pinto, H. P. de O., Raiman, J., … Zhang, S. (2019). *Dota 2 with Large Scale Deep Reinforcement Learning*. http://arxiv.org/abs/1912.06680

Palazzo, A., & Brozović, N. (2014). The role of groundwater trading in spatial water management. *Agricultural Water Management*, *145*, 50–60. https://doi.org/10.1016/j.agwat.2014.03.004

Parkash, V., & Singh, S. (2020). A Review on Potential Plant-Based Water Stress Indicators for Vegetable Crops. *Sustainability 2020, Vol. 12, Page 3945*, *12*(10), 3945. https://doi.org/10.3390/SU12103945

Perkin, J. S., Starks, T. A., Pennock, C. A., Gido, K. B., Hopper, G. W., & Hedden, S. C. (2019). Extreme drought causes fish recruitment failure in a fragmented Great Plains riverscape. *Ecohydrology*, *12*(6). https://doi.org/10.1002/eco.2120

Poli, R. (2008). Analysis of the Publications on the Applications of Particle Swarm Optimisation. *Journal of Artificial Evolution and Applications*, *2008*, 1–10. https://doi.org/10.1155/2008/685175

Qian, B., Hayhoe, H., & Gameda, S. (2005). Evaluation of the stochastic weather generators LARS-WG and AAFC-WG for climate change impact studies. *Climate Research*, *29*(1), 3–21. https://doi.org/10.3354/cr029003

Rabault, J., Kuchta, M., Jensen, A., Réglade, U., & Cerardi, N. (2019). Artificial neural networks trained through deep reinforcement learning discover control strategies for active flow control. *Journal of Fluid Mechanics*, *865*, 281–302. https://doi.org/10.1017/JFM.2019.62

Raes, D., Steduto, P., Hsiao, T. C., & Fereres, E. (2009). AquaCrop - The FAO Crop Model to Simulate Yield Response to Water: II. Main Algorithms and Software Description. *Agronomy Journal*, *101*(3), 438–447. https://doi.org/10.2134/agronj2008.0140s

Ragab, R. (2002). A holistic generic integrated approach for irrigation, crop and field management: the SALTMED model. *Environmental Modelling & Software*, *17*(4), 345–361.

Rodriguez, & Ober. (2019). AquaCropR: Crop Growth Model for R. *Agronomy*, *9*(7), 378. https://doi.org/10.3390/agronomy9070378

Rosa, L., Chiarelli, D. D., Rulli, M. C., Dell'Angelo, J., & D'Odorico, P. (2020). Global

agricultural economic water scarcity. *Science Advances*, *6*(18). https://doi.org/10.1126/SCIADV.AAZ6031/SUPPL_FILE/AAZ6031_TABLES_S2_TO_S5.XLSX

Rosenzweig, C., Elliott, J., Deryng, D., Ruane, A. C., Müller, C., Arneth, A., Boote, K. J., Folberth, C., Glotter, M., Khabarov, N., Neumann, K., Piontek, F., Pugh, T. A. M., Schmid, E., Stehfest, E., Yang, H., & Jones, J. W. (2014). Assessing agricultural risks of climate change in the 21st century in a global gridded crop model intercomparison. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(9), 3268–3273. https://doi.org/10.1073/pnas.1222463110

Salman, M., García-Vila, M., Fereres, E. (Elias), Raes, D., & Steduto, P. (2021). The AquaCrop model – Enhancing crop water productivity. In *The AquaCrop model – Enhancing crop water productivity*. FAO. https://doi.org/10.4060/cb7392en

Sammis, T. W. (1980). Comparison of Sprinkler, Trickle, Subsurface, and Furrow Irrigation Methods for Row Crops. *Agronomy Journal*, *72*(5), 701–704. https://doi.org/10.2134/agronj1980.00021962007200050002x

Sandhu, R., & Irmak, S. (2019). Performance of AquaCrop model in simulating maize growth, yield, and evapotranspiration under rainfed, limited and full irrigation. *Agricultural Water Management*, *223*, 105687. https://doi.org/10.1016/j.agwat.2019.105687

Saxton, K. E., & Rawls, W. J. (2006). Soil Water Characteristic Estimates by Texture and Organic Matter for Hydrologic Solutions. *Soil Science Society of America Journal*, *70*(5), 1569–1578. https://doi.org/10.2136/sssaj2005.0117

Scanlon, B. R., Faunt, C. C., Longuevergne, L., Reedy, R. C., Alley, W. M., McGuire, V. L., & McMahon, P. B. (2012). Groundwater depletion and sustainability of irrigation in the US High Plains and Central Valley. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(24), 9320–9325. https://doi.org/10.1073/pnas.1200311109

Scheierling, S. M., Loomis, J. B., & Young, R. A. (2006). Irrigation water demand: A meta-analysis of price elasticities. *Water Resources Research*, *42*(1). https://doi.org/10.1029/2005WR004009

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Openai, O. K. (2017). *Proximal Policy Optimization Algorithms*. https://doi.org/10.48550/arxiv.1707.06347

Schütze, N., Kloss, S., Lennartz, F., Al Bakri, A., & Schmitz, G. H. (2012). Optimal planning and operation of irrigation systems under water resource constraints in Oman considering climatic uncertainty. *Environmental Earth Sciences*, *65*(5), 1511–1521. https://doi.org/10.1007/s12665-011-1135-4

Scott, C. A., Vicuña, S., Blanco-Gutiérrez, I., Meza, F., & Varela-Ortega, C. (2014). Irrigation efficiency and water-policy implications for river basin resilience. *Hydrology and Earth System Sciences*, *18*(4), 1339–1348. https://doi.org/10.5194/HESS-18-1339-2014

Semenov, M. A. (2009). Impacts of climate change on wheat in England and Wales. *Journal of The Royal Society Interface*, *6*(33), 343–350. https://doi.org/10.1098/rsif.2008.0285

Semenov, M. A., & Barrow, E. M. (1997). Use of a stochastic weather generator in the development of climate change scenarios. *Climatic Change*, *35*(4), 397–414. https://doi.org/10.1023/A:1005342632279

Sentelhas, P. C., Gillespie, T. J., & Santos, E. A. (2010). Evaluation of FAO Penman–Monteith and alternative methods for estimating reference evapotranspiration with missing data in Southern Ontario, Canada. *Agricultural Water Management*, *97*(5), 635–644. https://doi.org/10.1016/j.agwat.2009.12.001

Sha, J., Li, X., & Wang, Z. L. (2019). Estimation of future climate change in cold weather areas with the LARS-WG model under CMIP5 scenarios. *Theoretical and Applied Climatology*, *137*(3–4), 3027–3039. https://doi.org/10.1007/s00704-019-02781-4

Sharma, K., Irmak, S., & Kukal, M. S. (2021). Propagation of soil moisture sensing uncertainty into estimation of total soil water, evapotranspiration and irrigation decision-making. *Agricultural Water Management*, *243*, 106454. https://doi.org/10.1016/j.agwat.2020.106454

Siebert, S., Kummu, M., Porkka, M., Döll, P., Ramankutty, N., & Scanlon, B. R. (2015). A global data set of the extent of irrigated land from 1900 to 2005. *Hydrology and Earth System Sciences*, *19*(3), 1521–1545. https://doi.org/10.5194/HESS-19-1521-2015

Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., Lillicrap, T., Simonyan, K., & Hassabis, D. (2018). A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, *362*(6419), 1140–1144. https://doi.org/10.1126/science.aar6404

Soulis, K. X., Elmaloglou, S., & Dercas, N. (2015). Investigating the effects of soil moisture sensors positioning and accuracy on soil moisture based drip irrigation scheduling systems. *Agricultural Water Management*, *148*, 258–268. https://doi.org/10.1016/j.agwat.2014.10.015

Steduto, P., Hsiao, T. C., Raes, D., & Fereres, E. (2009). AquaCrop—The FAO Crop Model to Simulate Yield Response to Water: I. Concepts and Underlying Principles. *Agronomy Journal*, *101*(3), 426–437. https://doi.org/10.2134/agronj2008.0139s

Suchato, P., Mieno, T., Schoengold, K., & Foster, T. (2022). The potential for moral hazard behavior in irrigation decisions under crop insurance. *Agricultural Economics*, *53*(2), 257–273. https://doi.org/10.1111/AGEC.12676

Sun, L., Yang, Y., Hu, J., Porter, D., Marek, T., & Hillyer, C. (2017). Reinforcement Learning Control for Water-Efficient Agricultural Irrigation. *2017 IEEE International Symposium on Parallel and Distributed Processing with Applications and 2017 IEEE International Conference on Ubiquitous Computing and Communications (ISPA/IUCC)*, 1334–1341. https://doi.org/10.1109/ISPA/IUCC.2017.00203

Sutton, R. S., Barto, A. G., & others. (1998). *Introduction to reinforcement learning* (Vol. 135). MIT press Cambridge.

Taghvaeian, S., Andales, A. A., Allen, L. N., Kisekka, I., O'Shaughnessy, S. A., Porter, D. O., Sui, R., Irmak, S., Fulton, A., & Aguilar, J. (2020). Irrigation Scheduling for Agriculture in the United States: The Progress Made and the Path Forward. *Transactions of the ASABE*, *63*(5), 1603–1618. https://doi.org/10.13031/trans.14110

Tenreiro, T. R., García-Vila, M., Gómez, J. A., Jiménez-Berni, J. A., & Fereres, E. (2021). Using NDVI for the assessment of canopy cover in agricultural crops within modelling research. *Computers and Electronics in Agriculture*, *182*, 106038. https://doi.org/10.1016/J.COMPAG.2021.106038

Turral, H., Svendsen, M., & Faures, J. M. (2010). Investing in irrigation: Reviewing the past and looking to the future. *Agricultural Water Management*, *97*(4), 551–560. https://doi.org/10.1016/J.AGWAT.2009.07.012

UNL. (2019). *Soil Water Content- and Soil Matric Potential-Based Irrigation Trigger Values for Different Soil Types*. https://cropwatch.unl.edu/2019/SWC-SMP-irrigation-trigger-values

URNRD. (2019). *Groundwater Management*. https://www.urnrd.org/programs-regulations/groundwater-management

USDA-NASS. (2018). *2018 Irrigation and Water Management Survey*. *3*(November). https://www.nass.usda.gov/Publications/AgCensus/2017/Online_Resources/Farm_and_Ranch_Irrigation_Survey/fris.pdf

USDA. (2019). *National Agricultural Statistics Service - Charts and Maps - Prices Received: Corn Prices Received by Month, US*. https://www.nass.usda.gov/Charts_and_Maps/Agricultural_Prices/pricecn.php

USDA. (2022). *Crop Insurance For American Farmers and Ranchers*. https://www.rma.usda.gov/

Vadez, V., Kholova, J., Zaman-Allah, M., Belko, N., Vadez, V., Kholova, J., Zaman-Allah, M., & Belko, N. (2013). Water: the most important 'molecular' component of water stress tolerance research. *Functional Plant Biology*, *40*(12), 1310–1322. https://doi.org/10.1071/FP13149

Vaishali, S., Suraj, S., Vignesh, G., Dhivya, S., & Udhayakumar, S. (2018). Mobile integrated smart irrigation management and monitoring system using IOT. *Proceedings of the 2017 IEEE International Conference on Communication and Signal Processing, ICCSP 2017*, *2018-January*, 2164–2167. https://doi.org/10.1109/ICCSP.2017.8286792

Van der Kooij, S., Zwarteveen, M., Boesveld, H., & Kuper, M. (2013). The efficiency of drip irrigation unpacked. *Agricultural Water Management*, *123*, 103–110. https://doi.org/10.1016/J.AGWAT.2013.03.014

Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H.,

Powell, R., Ewalds, T., Georgiev, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I., Huang, A., Sifre, L., Cai, T., Agapiou, J. P., Jaderberg, M., … Silver, D. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, *575*(7782), 350–354. https://doi.org/10.1038/s41586-019-1724-z

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., … Vázquez-Baeza, Y. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, *17*(3), 261–272. https://doi.org/10.1038/s41592-019-0686-2

Wang, D., & Cai, X. (2009). Irrigation Scheduling—Role of Weather Forecasting and Farmers' Behavior. *Journal of Water Resources Planning and Management*, *135*(5), 364–372. https://doi.org/10.1061/(ASCE)0733-9496(2009)135:5(364)

Water, U. (2017). *Value of Using Sensors to Manage Irrigation and Tips for Proper Installation*. https://water.unl.edu/article/agricultural-irrigation/value-using-sensors-manage-irrigation-and-tips-proper-installation

Wichelns, D. (2010). *Agricultural Water Pricing: United States*. https://doi.org/10.1787/786736626756

Wilhelmi, O. V., & Wilhite, D. A. (2002). Assessing vulnerability to agricultural drought: A Nebraska case study. *Natural Hazards*, *25*(1), 37–58. https://doi.org/10.1023/A:1013388814894

Wisser, D., Frolking, S., Douglas, E. M., Fekete, B. M., Vörösmarty, C. J., & Schumann, A. H. (2008). Global irrigation water demand: Variability and uncertainties arising from agricultural and climate data sets. *Geophysical Research Letters*, *35*(24). https://doi.org/10.1029/2008GL035296

Woltering, L., Ibrahim, A., Pasternak, D., & Ndjeunga, J. (2011). The economics of low pressure drip irrigation and hand watering for vegetable production in the Sahel. *Agricultural Water Management*, *99*(1), 67–73. https://doi.org/10.1016/J.AGWAT.2011.07.017

Yang, H. S., Dobermann, A., Lindquist, J. L., Walters, D. T., Arkebauer, T. J., & Cassman, K. G.

(2004). Hybrid-maize - A maize simulation model that combines two crop modeling approaches. *Field Crops Research*, *87*(2–3), 131–154. https://doi.org/10.1016/j.fcr.2003.10.003

Yang, Y., Hu, J., Porter, D., Marek, T., Heflin, K., & Kong, H. (2020). Deep Reinforcement Learning-Based Irrigation Scheduling. *Transactions of the ASABE*, *63*(3), 549–556. https://doi.org/10.13031/trans.13633

Yinglan, A., Wang, G., Hu, P., Lai, X., Xue, B., & Fang, Q. (2022). Root-zone soil moisture estimation based on remote sensing data and deep learning. *Environmental Research*, *212*, 113278. https://doi.org/10.1016/J.ENVRES.2022.113278

Young, R., Foster, T., Mieno, T., Valocchi, A., & Brozović, N. (2021). Hydrologic-Economic Trade-offs in Groundwater Allocation Policy Design. *Water Resources Research*, *57*(1), e2020WR027941. https://doi.org/10.1029/2020WR027941

Zellner, M., García, G. A., Bert, F., Massey, D., & Nosetto, M. (2020). Exploring reciprocal interactions between groundwater and land cover decisions in flat agricultural areas and variable climate. *Environmental Modelling and Software*, *126*, 104641. https://doi.org/10.1016/j.envsoft.2020.104641

Zhang, J., Guan, K., Peng, B., Jiang, C., Zhou, W., Yang, Y., Pan, M., Franz, T. E., Heeren, D. M., Rudnick, D. R., Abimbola, O., Kimm, H., Caylor, K., Good, S., Khanna, M., Gates, J., & Cai, Y. (2021). Challenges and opportunities in precision irrigation decision-support systems for center pivots. *Environmental Research Letters*, *16*(5), 053003. https://doi.org/10.1088/1748-9326/ABE436

Zhang, L., Wang, H., Chen, Z., Fang, Z., Sun, D., & Liu, G. (2020). Empirical Analysis and Countermeasures of the Irrigation Efficiency Paradox in the Shenwu Irrigation Area, China. *Water 2020, Vol. 12, Page 3142*, *12*(11), 3142. https://doi.org/10.3390/W12113142

Zhao, W., Queralta, J. P., & Westerlund, T. (2020). Sim-to-Real Transfer in Deep Reinforcement Learning for Robotics: a Survey. *2020 IEEE Symposium Series on Computational Intelligence, SSCI 2020*, 737–744. http://arxiv.org/abs/2009.13303

Zheng, H., Bian, Q., Yin, Y., Ying, H., Yang, Q., & Cui, Z. (2018). Closing water productivity

gaps to achieve food and water security for a global maize supply. *Scientific Reports*, *8*(1), 14762. https://doi.org/10.1038/s41598-018-32964-4

Zhu, C., Kobayashi, K., Loladze, I., Zhu, J., Jiang, Q., Xu, X., Liu, G., Seneweera, S., Ebi, K. L., Drewnowski, A., Fukagawa, N. K., & Ziska, L. H. (2018). Carbon dioxide (CO2) levels this century will alter the protein, micronutrients, and vitamin content of rice grains with potential health consequences for the poorest rice-dependent countries. *Science Advances*, *4*(5). https://doi.org/10.1126/SCIADV.AAQ1012/SUPPL_FILE/AAQ1012_SM.PDF

Zotarelli, L., Scholberg, J. M., Dukes, M. D., Muñoz-Carpena, R., & Icerman, J. (2009). Tomato yield, biomass accumulation, root distribution and irrigation water use efficiency on a sandy soil, as affected by nitrogen rate and irrigation scheduling. *Agricultural Water Management*, *96*(1), 23–34. https://doi.org/10.1016/J.AGWAT.2008.06.007