Social Choice and Welfare
https://doi.org/10.1007/s00355-023-01451-9

ORIGINAL PAPER

Monetizing the externalities of animal agriculture: insights from an inclusive welfare function

Kevin Kuruc[1,2,3] · **Jonathan McFadden**[2]

Received: 26 June 2020 / Accepted: 21 January 2023

Abstract

Animal agriculture encompasses global markets with large externalities from animal welfare and greenhouse gas emissions. We formally study these social costs by embedding an animal inclusive social welfare function into a climate-economy model that includes an agricultural sector. The total external costs are found to be large under the baseline parameterization. These results are driven by animal welfare costs, which themselves are due to an assumption that animal lives are worse than nonexistence. Though untestable—and perhaps controversial—we find support for this qualitative assumption and demonstrate that our results are robust to a wide range of its quantitative interpretations. Surprisingly, the environmental costs play a comparatively small role, even in sensitivity analyses that depart substantially from our baseline case. For the model to find that beef, a climate-intensive product, has a larger total externality than poultry, an animal-intensive product, we must simultaneously reduce the animal welfare externality to 1% of its baseline level and increase climate damages roughly 35-fold. Correspondingly, the model implies both that the animal agriculture sector is much larger than its optimal level and that considerations across products ought to be dominated by animal welfare, rather than climate, effects.

Kevin Kuruc and Jonathan McFadden contributed equally to this work.

✉ Kevin Kuruc
kkuruc@ou.edu

Jonathan McFadden
jmcfadden@ou.edu

1 Population Wellbeing Initiative, University of Texas at Austin, 305 E 23rd St., Austin 78712, TX, USA

2 Department of Economics, University of Oklahoma, 308 Cate Center Drive, Norman 73072, OK, USA

3 Global Priorities Institute, University of Oxford, Trajan House, Mill Street, Oxford OX2 0DJ, UK

Published online: 16 March 2023

1 Introduction

Raising animals for human consumption creates a host of important issues. Among the most pressing are the large effects on participants not represented in the current market: both the animals themselves and future humans, through the greenhouse gas (GHG) emissions of this sector, have stakes in this market. These—potentially large—unpriced externalities imply that current global production in animal agriculture may be meaningfully different than its efficient level.

This paper develops and applies a unified framework to jointly measure the animal welfare and climate externalities of animal agriculture. Within the disciplines of ethics and climate science, these issues respectively comprise large sub-disciplines. Notably, however, there are few attempts to quantify either in a welfarist framework that allows for economic measurement, i.e., a dollar value society should be willing to pay to reduce them. Indeed, we know of no work that attempts to monetize the animal welfare externality of industrial farming, let alone in a unified framework that allows for straightforward addition and comparison with the better-studied climate costs. This paper performs this exercise in a fully specified economic-welfarist framework and finds these costs to be jointly substantial with the possibility of very large animal welfare costs, depending on assumptions regarding the quality of animal lives.

We begin by formalizing a population-sensitive, animal-inclusive, social welfare function (SWF). The research question concerns counterfactually unborn animals, which requires stances not only on human-animal comparisons, but also on the social value of new existences. Following numerous ethicists and economists, we work with a generalized totalist utilitarian welfare function (Blackorby et al. 1995): social welfare is the (possibly weighted) sum of utility across all beings and time. While the field of population ethics—the study of how to rank social outcomes with different population sizes—lacks consensus over this issue (Arrhenius 2000), our choice has attractive features even for those who do not share this totalist view. First, we deduce that it produces a lower-bound on the welfare costs of animal production relative to a broad class of alternative SWFs. This is both because strict utilitarianism ignores sources of potential harms beyond the creation and discontinuation of streams of experiences in this market (Korsgaard 2018) and because totalist population criteria will be the friendliest towards adding lives that are not terribly good. Additionally, this function permits a simple analytical representation of the marginal cost of an animal product: the sum of all (discounted utility) costs to future humans plus the lifetime utility of the animal to be used for human consumption. The former has a natural analog in the social cost of carbon; it is conceptually simple to extend these costs to a different GHG-producing activity. The latter is less understood and requires novel quantitative stances on ethical parameters.

The most consequential of these ethical parameters regards how the life of an industrially ("factory") farmed animal compares to non-existence.[1] Using an analogous welfare function, Espinosa and Treich (2021) show that it is welfare-enhancing to reduce the size of this sector *if and only if* farmed animals do not have a "life worth living." In our model, because of the additional climate externality which is always negative, this becomes a sufficient (but not strictly necessary) condition of the sector generating a negative welfare externality. Appealing to various lines of reasoning, we argue that these animals likely do not have such worthwhile lives, implying that the existence externality is negative in our baseline analysis.[2] This stance is not universally accepted (Tännsjö 2016; Thompson 2020), and so the baseline results are most accurately viewed as an exploration of these costs *if* farmed animals do not have lives worth living. Throughout the paper we discuss what can be learned from our model under alternative views on farmed animal lives.

This animal-inclusive welfare function is applied within an economic model, DICE-FARM, built in companion research that examines the climate costs of a range of dietary choices (Errickson et al. 2021). DICE-FARM is a modification of the Dynamic Integrated Climate-Economy (DICE) model (Nordhaus 1992, 2017) that includes an animal agriculture sector. Formally, Errickson et al. (2021) extends DICE to capture the effects of non-CO_2 GHGs from livestock production—methane (CH_4) and nitrous oxide (N_2O)—and includes a sector that generates these emissions as a by-product of animal production. In this paper, aside from generalizing the welfare function, we further enrich that model by accounting for the number of animal life years used to produce meat within the farm sector.

Our baseline results imply that the annual social costs of an average American diet are very large: the total welfare costs to future humans and animals are monetized at values on the $100,000 order of magnitude per diet, per year. In other words, we estimate that the social loss generated from one individual's annual meat consumption is greater than $100,000. The animal welfare externality accounts for nearly all of this large sum, underscoring the degree to which the results rely on difficult-to-test assumptions on the parameter governing the quality of farmed animals' lives. However, in robustness analyses, we demonstrate that the animal welfare externality remains the dominant consideration under nearly any quantitative representation of the qualitative claim that industrially farmed animals have lives not worth living. Under the opposing claim that these animals do in fact have net-pleasurable existences, the results are quickly reversed. Though, due to the lower-bound nature of the totalist utilitarian framework in this application, we are less confident about the implications of our results under an assumption of net-pleasurable lives. In either case, at both the annual diet and individual serving level, chickens are

[1] The physical and mental well-being of livestock vary significantly by production method and country. We assume that an industrially farmed animal is representative given the dominance of this production, not only in the United States, but in other middle- and high-income countries (Gerber et al. 2013).

[2] We apply the term "existence externality" or "existence cost" in a manner distinct from the use of "existence value" in environmental economics (Krutilla 1967). There, humans derive utility from a natural resource or living organism based on knowledge of its existence. In our setting, existence cost is taken to mean the social cost realized from livestock that experience net negative utility from living.

the main source of these welfare effects. Poultry has become a dominant source of protein, and the meat produced per bird is low, so each serving requires more animal lives; as a result, more than 68 billion birds are raised and slaughtered annually (FAO 2021b).

There are two straightforward policy implications arising from the baseline analysis that we proceed to study. First, the large externalities associated with this market suggest that optimal policy would reduce its size. To formalize this, we add to the model private production costs and utility benefits received by consumers. Following directly from the large external costs, the welfare maximizing level of animal agriculture is much smaller than the unregulated outcome. The functional form we choose for utility implies steeply increasing marginal benefits as meat consumption declines, so the sector is not fully eliminated in our calibration, though we do not wish to stress the exact quantitative results given the many uncertainties underlying this exercise. Rather, the results suggest that any modest policy proposals to reduce the sector will be supported in our framework.

The second, and arguably more decision relevant implication, is that consuming poultry instead of beef has the possibility of being welfare reducing despite the well-known climate benefits of this substitution. A further product-level optimization problem inheriting our baseline assumptions would merely reflect that poultry has higher social costs than beef and prioritize poultry reductions. Instead, we study the robustness of the relative marginal social costs between these products to assumptions underpinning the climate and animal welfare costs. This allows us to directly study the parameter combinations that can rationalize a substitution from beef to poultry on welfarist grounds. We find that only a small corner of the relevant parameter space is able to support such a substitution in this framework. For example, one such combination that equates the externalities of beef and poultry implies that our baseline analysis overestimates the animal welfare externality 100-fold *and* that DICE-FARM underestimates climate costs by roughly 35-fold; other parameters that reverse the original result are similarly extreme relative to our baseline. Given the unique uncertainties in this context, we acknowledge it is plausible that our baseline is in error by these magnitudes, though it is informative to recognize that *both* dimensions would need to be seriously modified to generate qualitatively different implications.

This paper contributes first and foremost to the sub-field of environmental economics concerned with animal welfare. A recent landmark in this literature is Norwood and Lusk (2011) which serves as a thorough treatment of the economics and ethics of farmed animal welfare. The key difference in our paper is that we incorporate farmed animals into the welfare function directly—their well-being matters for their sake—whereas Norwood and Lusk (2011) mainly ask how humans value animal welfare (see also Fleurbaey and Van der Linden 2021). The subset of papers drawing on the inclusive concept of welfare that we apply is even smaller; Johansson-Stenman (2018) and Carlier and Treich (2020) highlight this missing literature and call for animal welfare to be directly included in economic analysis. Very few papers have attempted this. Blackorby and Donaldson (1992) and Espinosa and Treich (2021) study the properties of a joint-maximization problem over the quantity and quality of animal lives in settings where humans choose how many animals live.

Within the context of climate change, Hsiung and Sunstein (2006) and Budolfson and Spears (2019) find that inclusion of wild animal welfare greatly matters for valuing climate damages, thus altering the trajectory of optimal policy. In a similar spirit, Ng (1995) and Groff and Ng (2019) study baseline wild animal welfare through an evolutionary economic framework. We contribute the first estimate of (inclusive) welfare effects coming from the industrial farming sector.

To do so, we draw on and further contribute to three distinct areas of study. First, within the field of applied ethics, much has been written generally on animal welfare. Tracing its (Western) history to Bentham (1789), modern work on this topic was catalyzed by Singer (1975). Despite the well-known utilitarian frameworks of Bentham (1789) and Singer (1975), there is wide agreement from researchers across ethical frameworks that animal welfare and animal rights deserve more consideration than they receive at present (Kagan 2019; Korsgaard 2018). It is in response to this ethical consensus that Carlier and Treich (2020) make their call for economists to take on the challenge of evaluating animal welfare in policy decisions where it is likely to be affected.

Second, since our study concerns the number of farmed animals who ever exist, we contribute to the applied literature on social choice regarding variable population sizes. Despite the challenges of making social choices over populations of non-constant size (Arrhenius 2000; Blackorby et al. 2005; Greaves 2017), in any plausible conceptualization these considerations will likely be quantitatively important if included (see Lawson and Spears (2020) for an example in human settings). We draw on the field of population ethics which has proposed a variety of rigorous frameworks for dealing with choices of this type. Although universal consensus does not exist regarding the proper social welfare function, we follow a tradition in economics that employs a generalized critical-level total utilitarian approach (e.g. Blackorby et al. 1995). This choice is in line with the small set of aforementioned papers that have formalized inter-species population ethics comparisons (Blackorby and Donaldson 1992; Espinosa and Treich 2021; Budolfson and Spears 2019).

Third, as the main purpose of the paper is to sum welfare costs and compare them to widely-recognized externalities, we draw on results and models from environmental economics.[3] We direct the reader to companion research on the relationship between livestock production and climate change that develops and analyzes the integrated assessment model on which this paper relies (Errickson et al. 2021). The methods employed therein closely mirror measurement of the social cost of carbon (e.g. Nordhaus 2017). Outside of Errickson et al. (2021), we know of only a few papers that make attempts to price the GHG externality of animal agriculture (Wirsenius et al. 2011; Springmann et al. 2016).

[3] It must be noted that, apart from substantial climate change impacts (Ripple et al. 2014), livestock farming results in sizeable land use change (Foley et al. 2005), biodiversity loss (Newbold et al. 2015), air and water pollution (Diaz and Rosenberg 2008), drawdown of freshwater sources (Wada et al. 2010), and antimicrobial resistence (Innes et al. 2020). See Errickson et al. (2021) for a more extensive discussion of these latter impacts, which we do not consider in this work.

The paper proceeds as follows. Section 2 formalizes the welfare function and the ethical assumptions we employ throughout the paper. Section 3 presents the quantitative exercises measuring the joint and independent externalities in this sector. Section 4 discusses the policy implications of these findings and studies their robustness. Section 5 concludes.

2 Inclusive welfare and the social costs of animal farming

In this section we present the welfare function to be used throughout the paper and provide arguments for quantifying key ethical parameters. First, we formalize the total utilitarian framework and derive the marginal welfare cost of raising animals for human consumption. We then discuss how we parameterize the ethical components introduced by this welfare function. The section concludes with a discussion of the relation between our framework and other leading alternatives.

2.1 Interspecies critical-level utilitarianism

The assumed welfare function follows the generalized critical-level utilitarian framework of Blackorby et al. (1995). We extend the function to include animal well-being as in Blackorby and Donaldson (1992) and Espinosa and Treich (2021). Social welfare is defined as the sum of (socially discounted) intra-period welfare:

$$SWF = \sum_{t=t_0}^{\infty} \left(\frac{1}{1+\rho}\right)^t \overbrace{\Bigg[\underbrace{\sum_{i=1}^{P_t} (u_{i,t} - \underline{u})}_{\text{Human Utility}} + \theta \underbrace{\sum_{j=1}^{A_t} (u_{j,t}^A - \underline{u})}_{\text{Animal Utility}} \Bigg]}^{\text{Total Within Period Utility}}. \tag{1}$$

Within periods the function sums human and animal populations' (P_t, A_t, respectively) welfare above some utility threshold $\underline{u}$, which we take to represent a level that is, from the agent's subjective perspective, net-neutral relative to non-existence. Note that by assuming the critical level to be a subjectively neutral existence our generalized set up becomes the special case of total utilitarianism. In what we are referring to as a neutral life, the intensity weighted duration of pleasurable experiences from the point of view of the individual equals the intensity weighted duration of displeasurable experiences; implicitly this hedonistic framing follows the tradition of Bentham (1789).

While holding fixed the concept of a neutral life across species, we acknowledge that the same external activities and stimuli produce very different internal experiences across species. An hour rolling about in mud may create a net-positive experience for pigs while creating a near-neutral (or negative) experience for humans. All of the uncertainty regarding these subjective differences across agents is subsumed in the utility terms. We note that this is not unlike the treatment of utility differences across humans in standard settings; if two humans value different experiences with different intensities, we would find it natural for those features to be part of their respective utility functions.

Likewise following from our hedonistic conception of utility, we set the welfare weight on animal utility, θ, equal to one. This choice follows a long anti-speciesist tradition (Singer 1975). To reiterate, we recognize the tremendous uncertainty regarding how external states and stimuli translate to subjective experiences across species. One particularly relevant concern here is over the range and depth of potential utilities. It may indeed be the case that human brains generate more extreme subjective experiences, say, if contemplating a beautiful piece of art is a better human experience than the best experience a pig can have, or conversely, if losing a loved one produces more grief in humans than pigs can experience. Setting θ to one only requires that *equivalent* subjective experiences are valued equally; we make no claims that other species can, in fact, have similarly good (bad) lives if things go well (poorly) for them. Unlike the constant critical levels across species, where we merely shift experiential differences into utilities rather than species-specific critical levels, the assumption of θ equal to one is ethically substantive. To see this, consider a scenario where a human and cow are in equal (subjective) pain and there is a single pain killer available. If one believes it would be morally preferable to alleviate the pain of the human, this is an implicit endorsement of $\theta < 1$. Without resolving such difficult dilemmas, we note throughout the presentation of results how they depend on θ.

The additively separable structure of this welfare function implies an additional normative stance that we explicitly note before proceeding: adding a net-pleasurable life increases social welfare. Likewise, adding a net-displeasurable life reduces social welfare. These implications are first order concerns when studying a market in which humans directly control the number of animals who exist. We recognize that these implications, particularly the former, strike many as non-obvious. It is a longstanding challenge in the philosophical-economics literature to build a variable population welfare function that is consistent with a set of widely held intuitions (Arrhenius 2000; Greaves 2017). Totalist population criteria satisfy the most reasonable set of normative axioms in our view and the choice follows many economic studies concerned with variable population problems (Blackorby and Donaldson 1992; Espinosa and Treich 2021; Méjean et al. 2020). However, we note at the close of this section that this framing produces a lower-bound on the welfare costs of bringing animals into existence. Under competing conceptions of welfare, the high costs we estimate in Sect. 3.2 would be in fact higher. Normative disagreements on this issue need not lead to substantive disagreements with the analysis, results, and implications.

2.2 The marginal welfare costs of rearing farmed animals

A simple analytical expression for the social cost (or benefit) of raising an additional animal for consumption arises from this welfare function. For expositional simplicity, we temporarily assume all farmed animals have equal (annualized) utility, u^A, and introduce a new term, LS^A, to represent the lifespan of these animals. In our subsequent application we generalize the model to allow for differences in utility and lifespan across animals.

$$\frac{\partial SWF}{\partial A_0} = \left[\sum_t \left(\frac{1}{1+\rho} \right)^t \left(\sum_{i=1}^{P_t} \frac{\partial u_{i,t}}{\partial A_0} \right) \right] + \theta \times LS^A \times [u^A - \underline{u}] \tag{2}$$

The first term corresponds to the effects on humans, through time, of raising an additional animal today. We restrict these costs to the climate effects of farmed animals, ignoring local environmental and other externalities imposed by these operations. With this simplification, the human-cost term can be expanded as follows:

$$\frac{\partial u_{i,t}}{\partial A_0} = \frac{\partial u_{i,t}}{\partial T_t} \times \frac{\partial T_t}{\partial E_0} \times \frac{\partial E_0}{\partial A_0}. \tag{3}$$

Welfare lost to each individual is the product of three terms: (i) utility changes for person i in year t from a warming planet, $\frac{\partial u_{i,t}}{\partial T_t}$, (ii) temperature changes in year t from an additional unit of emissions today, $\frac{\partial T_t}{\partial E_0}$, and (iii) emission changes today from an additional animal raised today, $\frac{\partial E_0}{\partial A_0}$. To simplify notation, we denote GHG emissions as a scalar, E_0. In our application, E_0 is a three-dimensional vector that includes carbon dioxide, methane, and nitrous oxide.

The animal welfare term in Eq. (2) is the product of: (i) the difference in annualized utility from the critical level, $u^A - \underline{u}$, (ii) the lifespan of the animal, LS^A, and (iii) the welfare weight placed on animals, θ. The lifespan enters because the welfare function is defined per period (i.e., per year) and so u^A is implicitly defined as annualized utility.

Following simplifications in the climate-economics literature (Nordhaus 2017), we use global averages for both humans and animals and weight average utility by the population size in each period. For humans, the utility function is assumed to exhibit constant relative risk aversion (CRRA) and to depend only on per-capita consumption, $\bar{c}_t$:

$$\sum_{i=1}^{P_t} [u(\bar{c}_t) - \underline{u}] = P_t[u(\bar{c}_t) - \underline{u}] = P_t \left[\frac{\bar{c}_t^{1-\eta}}{1-\eta} - \underline{u} \right]. \tag{4}$$

This function is parameterized as in the DICE model and other climate-economy welfare calculations, save for $\underline{u}$ which we discuss in the following subsection. We postpone discussion of any utility garnered specifically from meat consumption to Sect. 4 where it becomes relevant for an optimal policy analysis.

2.3 Existence value and farmed animal welfare

The literature on the economics of animal welfare provides little guidance for calibrating the parameters that determine animal welfare costs. We begin by assuming that u^A is fixed over time. This simplification has no effect on our main results which study the marginal welfare cost of an additional animal today, as in Eq. (2).

We likewise assume that u^A is fixed across farmed animals. This is substantively important. Unfortunately, it is impossible to make confident statements about cross-species welfare at present. Rather than make such conjectures, we believe this uniform assumption has the benefit of rendering the analysis transparent. Where plausible cross-species differences in welfare would importantly influence the qualitative takeaways of our analysis we note this source of uncertainty. Choosing a magnitude for this now-fixed u^A—and more consequently its relation to $\underline{u}$—requires assumptions on unknowable quantities and experiences. We proceed with humility.

First, we follow studies that rely on similar calculations for humans and assume that $\underline{u}$ corresponds to lives lived somewhere near the international purchasing power parity (PPP) adjusted poverty line of \$1.90 per day (e.g. Tännsjö 2016; Méjean et al. 2020). That is, we assume life becomes better than non-existence once consumption levels are above the current internationally-defined poverty line. On the one hand, we recognize—as do others in this literature—it may seem demeaning to suggest that millions of human beings have lives that are not worth living. We acknowledge this concern.

Nevertheless, it is plausible that some human lives include more negative experiences than positive experiences. A shortcoming of our utility function is that it is only responsive to income, and so we are forced to represent the concept of a net-negative life by income levels despite the fact that many individuals below this threshold surely enjoy net-pleasurable lives. Conversely however, compelling arguments can be instead made that this \$1.90 threshold is too low. The additive nature of our welfare function leads to the well-known “Repugnant Conclusion” in which a world with an arbitrarily large population of lives just above $\underline{u}$ would be objectively better than the status quo (Parfit 1984). If it seems implausible that a world of many individuals living on \$1.91 per day (i.e., just barely net-positive on our calibration) is better than our current world, this suggests that our assumed critical level is too conservative.

To choose a value for u^A, the utility of a farmed animal, we draw on four lines of independent reasoning that lead us to set $u^A < \underline{u}$ in our baseline. First, there are many animals, both companion animals and those raised for food, for which we find it perfectly reasonable—even “humane”—to euthanize on the animal’s behalf. This implies that humans forecast net-negative experiences for these animals, else euthanasia would not be in the animal’s interest. The condition of a representative farmed animal is likely worse than living as a chronically ill or moribund house pet. By transitivity then, farmed animal lives would be below a neutral level (Matheny 2003).

Second, in the spirit of Pearce (2021) one could perform a (very imperfect) veil-of-ignorance exercise between living 1 year as a human at \$1.90 per day (equivalent to $\underline{u}$ by assumption) or living for 1 year as an industrially farmed animal. This exercise strains credulity as it requires not only imagining that one is being farmed, but having the subjective experience of an animal in these conditions. Nonetheless—if forced to choose under these tremendous uncertainties—we would expect the former to have more subjective happiness and fewer moments of stress, boredom, and pain. This implicitly suggests that we believe the utility of a farmed animal is, in expectation, lower than our assumed critical level.

Third, we can again leverage the "Repugnant Conclusion" for an informal proof by contradiction. Under our additive utilitarian welfare function, if factory farmed animals have net-pleasurable lives, some number of their existences can offset fewer (net-pleasurable) human lives. At the extreme then, a world of arbitrarily large factory farms and no human beings can be shown to generate more total well-being than our current world. To us, this seems an objectively worse state of affairs, and so it must be the case that industrially farmed animals do not contribute positively to social welfare. Equivalently, $u^A < \underline{u}$.

Finally, setting aside our imperfect personal reflections on this question, survey evidence suggests the dominant view is that industrially farmed animals do not have lives worth living. In Espinosa and Treich (2021) survey participants are described the conditions that broiler chickens—the most numerous farmed animal—experience in intensive indoor rearing practices. A large majority across students, philosophers, activists, and even farmers, view these lives as not worth living.[4] Additionally, while there exist dissenting voices contending that because these animals have needs provided (such as calories and shelter from predation) their lives are worth living (e.g. Tännsjö 2016; Thompson 2020), our reading of the philosophical literature on animal ethics suggests that it is the dominant view that the current state of industrial farming does not result in worthwhile lives. In accordance with arguments above, many believe that even lives with these minimal needs met may be negative on the whole; a human life of solitary confinement, for example, seems to plausibly fit this description.

Beyond these arguments remains quantitative difficulties that are impossible to resolve. As such, we set u^A at the arbitrary value of human-equivalent utility at \$1.00 per day, satisfying the condition that $u^A < \underline{u}$. We recognize that tremendous uncertainty surrounds this monetary choice, and even the broader statement that $u^A < \underline{u}$. Accordingly, we make explicit how our main results vary over a wide range of values for u^A.

2.4 Hedonistic total utilitarianism as a lower bound on welfare costs

The choice of hedonistic, totalist utilitarianism may strike some as inappropriate in light of several of its implications. For example, animals with vanishingly short lives—male chicks born at egg laying facilities, for example—are given no weight in our calculation. More broadly, the action of ending an otherwise worthwhile life is only represented as the lost opportunity of future utility for that being. Aside from this concern, our additive aggregation ignores other proposals within the social welfare literature to deal with the separate issues of inequality (e.g. Adler 2008; Zuber and Asheim 2012) and implications that arise from explicitly valuing population increases (Parfit 1984). Resolving these differences is beyond the scope of this paper. In what follows, however, we note that our choice will capture a lower-bound

[4] Participants were given a range of animal living conditions and asked to assess at which point lives became "worth living." Even for living conditions more pleasant than that of industrially farmed animals, most participants rated these as lives as not worth living.

on costs, and hence our ultimately large estimates serve as a starting point for tallying all possible negative welfare effects.

We first comment on the strictly hedonistic framing—that is, the only welfare loss resulting from raising an animal for food in this setting (other than the climate costs) is the animal's instantaneous suffering summed across its life. This implies that one pig living 1 year is equivalent to two pigs living six months each, despite two deaths occurring in the latter case. An important dimension of morality may be omitted by this, namely that animals may have some right not to be raised merely for slaughter. In this case, we could add costs within our framework to account for the act of killing as a violation of the animal's right to be an end in the Kantian sense. Doing so would clearly increase the total social costs of animal agriculture. Indeed, Korsgaard (2018) argues from such a Kantian framework that no animal agriculture is permissible if we grant animals moral status.

Regarding aggregative methods, we can divide competing theories between those with distributional concerns and those with concerns about large populations being socially desirable merely because they are large. Theories such as prioritarianism (Adler 2008), sufficientarianism (Shields 2012), and rank-discounted utilitarianism (Zuber and Asheim 2012) are proposals that give extra weight to the marginal utility of the worst off. As we assume farmed animal utilities are at or near the worst human experiences, welfare measures that prioritize the worst off are more sensitive to the plight of farmed animals than our utilitarian framework. Regarding concerns about large populations, the most widely cited alternative to totalist population criteria is instead averagist criteria which consider average welfare conditional on existence. Again, if farmed animals have lives near the bottom of the distribution of existences, an averagist welfare function will put substantial value on preventing their existence. These existences pull down average welfare more quickly than they pull down total welfare. This is most salient for lives only just not worth living—total welfare is negligibly impacted by the addition of such a life, but average welfare is pulled towards neutrality. As a consequence, frameworks that put some (or all) weight on averagist criteria will treat welfare costs as being at least as large as our totalist welfare function.

This discussion is not intended to refute competing theories nor defend total utilitarianism. Rather, it is a bounding exercise. Our results are at least as large as those that would obtain from the same economic exercise through any of the ethical frameworks discussed above.

3 Quantification of costs in an augmented IAM

We now describe the model used for the application—an augmented version of the DICE model—and present the results. We find that the welfare costs of global animal agriculture are very large in the case that animals do not have net-pleasurable existences: the monetized costs of producing the meat consumed for the Standard American Diet (SAD) for one person is on the order of $100,000 per year under our baseline parameters. In other words, eliminating the production of meat required for one individual's diet for 1 year confers social welfare benefits equal to the benefits of increasing annual global output by more than \$100,000. This is entirely driven by the negative existence value from the sheer volume of animals produced for food—notably chickens—and therefore varies with our choice of u^A. However, for nearly any value below neutrality, the welfare costs remain large. In the case that animals have worthwhile lives, the results are quickly reversed, highlighting the importance of this parameter for appropriate policy recommendations.

3.1 Model details: DICE-FARM with animal welfare

Our model builds on the DICE-FARM framework developed in Errickson et al. (2021). In that work, the focus is solely on pricing the climate externality from animal agriculture using a modified version of DICE (Nordhaus 2017), a leading integrated assessment model. The standard DICE model, like most macroeconomic IAMs, consist of four conceptual modules. The economic module uses current economic inputs to produce goods with a by-product of CO_2 emissions; an atmospheric module maps the history of emissions to the current stock of GHGs in the atmosphere; a climate module inputs the GHG stocks from the atmospheric module to compute temperature dynamics; and a damage module uses the temperature increases as negative inputs to the economic module. The basic trade-off is that output today increases utility directly but harms future utility indirectly through the climate-economy cycle. Section B of the online appendix contains model details.

Errickson et al. (2021) modifies DICE in two important ways to create DICE-FARM. First, an animal agricultural sector is included alongside the industrial-only output of the economic sector. This module produces meat for human consumption with the by-product of emissions from farmed animals. Emissions intensities, taken from life-cycle assessment analysis performed by the United Nations' Food and Agricultural Organization, reflect emissions from land use change, production of feedstuffs and other farming inputs, animal management, direct and indirect energy use, and post-farm activities (FAO 2021a). Because farmed animals contribute a quite different mix of potent GHGs—including methane and nitrous oxide—the climate module in DICE-FARM must also be modified to endogenize the evolution and impact of these gases. For this purpose, the FAIR climate module is used (Millar et al. 2017). This modification has the additional benefit of responding to requests to substitute the climate module within DICE for one that better reflects current scientific consensus (National Academies 2017). Additional details of these modifications and the resulting animal-environmental relationships can be found in Errickson et al. (2021).

For our welfarist exercises, we must further enrich DICE-FARM along two dimensions. First, the human-centric social welfare function in Errickson et al. (2021) is replaced by our animal-inclusive total utilitarian function. Second, we modify the farm sector to explicitly account for the number of animal-life years necessary to produce each unit of meat. For each type of meat, this is the product of (i) the number of animals slaughtered in a given year and (ii) average lifespans, which we (iii) divide by total global production. Table 1 summarizes animal life-years necessary to produce one serving size of 20 g of protein across the three animals.

Table 1 Life-years lived per 20 g protein

Meat product	Life-years
Beef	0.0017
Pork	0.0010
Chicken	0.0103

Life-years lived per serving of meat products, defined as 20 g protein (approximately one hamburger). For each product, life-years are computed as: number of animals slaughtered annually × lifespans of those animals ÷ servings produced annually

Table 2 Marginal external welfare costs of dietary decisions

	(1)	(2)	(3)	(4)
	Annual	20 g Protein		
	Non-vegetarian	Beef	Pork	Chicken
Total	122,836.64	56.31	32.79	325.35
Environmental	47.45	0.17	0.03	0.01
Animal welfare	122,789.19	56.14	32.76	325.34

External costs of dietary decisions in USD. Column (1) is computed by considering an extensive margin decision: whether or not to eat any meat, relative to the meat consumed in an average American diet. Columns (2)–(4) consider the intensive margin, for example the cost of eating one more hamburger. All values are computed using a discretized approximation of Equation (5)

Chickens require the most life-years per serving due to their much smaller body mass than pigs or cows.

With the inclusion of animal welfare in DICE-FARM, the two primary externalities of this sector have been accounted for, making the social costs of animal agriculture conceptually straightforward to compute.

3.2 Main results: animal inclusive social costs of meat eating

The main results of the paper are the social costs of various dietary decisions, D. We define and compute these social costs in a manner analogous to the social cost of carbon (Nordhaus 2017).

$$SC(\Delta D) = \frac{\frac{\Delta SWF}{\Delta D}}{\frac{\Delta SWF}{\Delta C}} \tag{5}$$

The numerator is the total welfare change associated with the dietary margin under consideration. We analyze both the extensive (vegetarian) margin at an annual time horizon and the intensive (per-meal) margin. The denominator is the welfare change from a marginal dollar of consumption. Conceptually, $SC(\Delta D)$ is the dollar change to which the dietary change is welfare-equivalent.

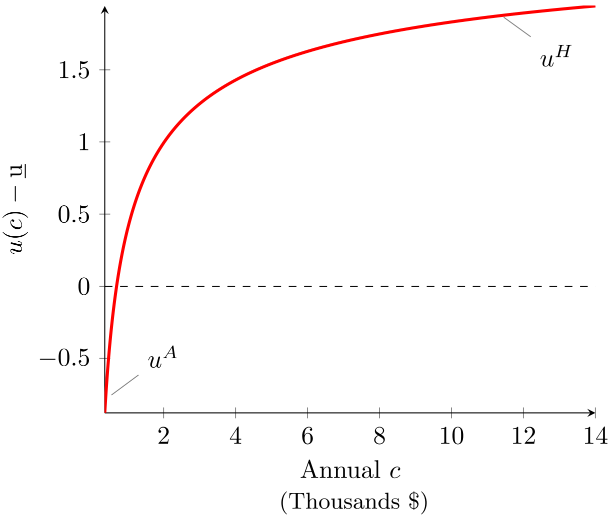

Fig. 1 Concavity implies low utility levels for animals. Per-capita utility as a function of income under baseline assumptions ($\eta = 1.45$). u^A is approximately −0.8 at $c = 0.365$ (thousands per year; equivalently 1 dollar per day). u^H is the utility derived at 2020 levels of global average consumption within the standard DICE model applied

The results of this exercise are striking under our baseline parameters,[5] depicted in Table 2. The cost of an annual non-vegetarian diet relative to a vegetarian diet is estimated to be \$122,837. Only \$47 of this cost comes from environmental factors. To be sure, \$47 of external climate costs per person, per year, are significant when summed across the many people consuming these products, and this may even be a lower-bound given that the DICE damage function is thought to underestimate climate damages (e.g. Weitzman 2012; Burke et al. 2015; Howard and Sterner 2017). However, we know of no adjustments to the climate module that can magnify these costs to the level of the baseline animal welfare costs (see Sect. 4.2). Bringing beings into existence with lives not worth living has significant welfare effects in this population-sensitive framework. Accordingly, the results when considered on the intensive margin (20 g of protein) are dominated by chicken despite the fact that a single serving imposes a mere penny's worth of environmental costs. The number of meals per chicken is small relative to other animals, and the value of each animal life-year is large in magnitude.

These large costs stem wholly from the assumption regarding farmed animals' deviation in utility from a neutral existence. The baseline results place their utility at the equivalent of a human life on \$1.00 per day. The concavity of the human utility function implies this is much worse than our assumed neutral existence in

[5] As stated earlier: the parametric welfare assumptions are: u^A is equal to the utility from a human life at \$1 per day, $\underline{u}$ is equal to human utility at \$1.90 per day, $\rho = 0.015$, and $\eta = 1.45$. The latter two parameters are taken directly from Nordhaus (2017).

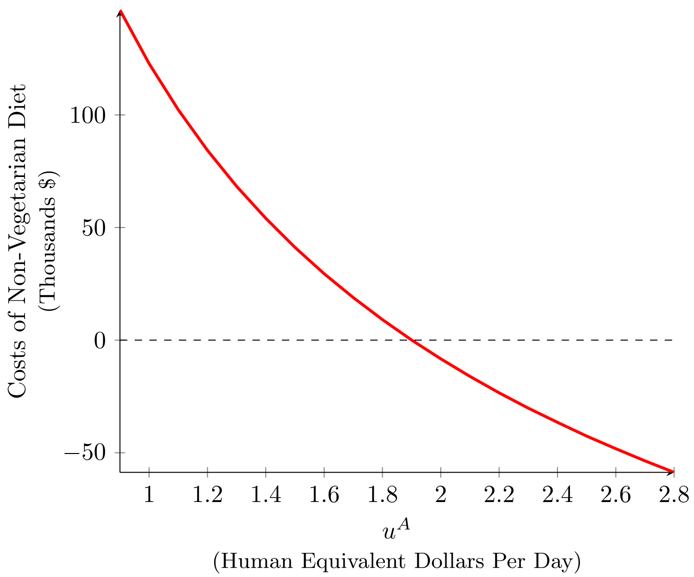

Fig. 2 Costs of non-vegetarian diet by farmed animal utility. Total annual costs of non-vegetarian diet, corresponding to Table 2 column (1), for different assumptions over farmed animal utility, u^A. The x-axis reports the human-equivalent-consumption that would generate the u^A used in that scenario, i.e, a value of \$1.50 here implies u^A is equivalent to the utility humans receive living at \$1.50 per day

a utility-sense. To see this, Fig. 1 plots $u(c) - \underline{u}$ for different values of c. Lives at \$1.00 and \$1.90 per day generate quite different levels of utility when this function is calibrated to standard values for the elasticity of marginal utility with respect to consumption.

The values in Table 2 reflect how much total global output would need to be increased to offset total social welfare losses resulting from adding some number of animal lives. Adding an animal life-year reduces total (inter-species) social welfare by about 0.8 utils. The global aggregate income gains necessary to increase human welfare by this same amount is large, approximately \$30,000, because the output gains are split among the world population. Each individual sees their consumption rise by a small quantity at a point where utility is relatively invariant to consumption. Thus, large total output gains are necessary to offset the utility losses associated with rearing an additional farmed animal (see Appendix A for a detailed derivation of this calculation).

Despite the natural ambiguity that arises in regards to the quality of life experienced by farmed animals—and the influence that this parameter has on the results—Fig. 2 shows that they remain large for nearly any choice in which animal lives are worse than neutral. This figure plots how annual social costs of an average diet varies based on different assumptions over u^A, expressed on the x-axis as human-consumption-equivalent welfare values. Our baseline paramaterization (that an animal life is the utility equivalent of living on \$1 per day as a human) corresponds to 1.0 on the x-axis where social costs are measured near \$125,000. As animal well-being increases, the social costs of raising them for food decrease. However, even as this value approaches neutral existence ($1.90), the costs remain well into the thousands of dollars.[6] One would have to be quite confident animal lives are nearly neutral in order for this to not be a first-order welfarist concern.

Alternatively, for the case in which farmed animal lives are worth living—values to the right of $1.90—the costs become nearly as large and *negative*. In such a case, our framework implies very large welfare *benefits* from producing farmed animals with sufficiently pleasurable lives. While important for demonstrating that the takeaways from the baseline analysis are immediately nullified when animals are assumed to have worthwhile lives, we put little stock in the additional implications on this domain. As noted in Sect. 2.4 our framework likely represents a lower-bound on costs, so a demonstration of negative costs provides little actionable guidance.

Figures 1 and 2 also highlight the importance of the curvature of the utility function. Optimal climate policy greatly depends on this curvature because it governs comparisons across agents of different wealth levels (Dasgupta 2008; LoPalo et al. 2019), as is implicit in our analysis. The online appendix shows that the large cost estimates are not driven by our assumptions over the elasticity of marginal utility (Fig. A1).

4 Policy implications: optimal sector size and product substitution

The baseline results in Sect. 3.2 point towards two distinct welfare-improving social choices. First, the large social costs of meat-eating in general suggest optimal levels of animal agriculture may be significantly lower than current levels. We formalize this conjecture by adding structure on the private consumption and production markets and performing an optimal policy exercise. Second, because the main estimates are driven by animal welfare considerations, substituting away from an animal-intensive meat (poultry) towards a climate-intensive meat (beef) appears welfare enhancing. In the case that cross-product substitution is a more active margin than uniform reductions, this is the more decision relevant finding and therefore deserves further exploration. There are significant uncertainties in both the animal welfare and climate aspects of the model that contribute to this finding; we therefore proceed by mapping the parameter space that supports the original result—and conversely, the parameter combinations required to overturn it. Relative to our baseline calibration, large modifications are necessary to generate the result that beef is more socially costly than poultry.

4.1 Optimal levels of animal production

Formalizing statements regarding optimal animal agriculture requires introducing utility benefits to humans from meat consumption. In a market equilibrium, private marginal benefits are equal to private marginal cost, and by documenting large

[6] At $1.90, the social cost is only the environmental cost, and thus not zero. These are indistinguishable on this graph due to the required scaling of the *y*-axis.

external costs at current levels of consumption, we can be confident that the market outcome is higher than its efficient level.[7] As the size of this sector decreases, however, the marginal benefits of meat consumption are likely to increase. In order to estimate at what point marginal private benefits equal marginal social cost, we must impose structure on the (human) utility function.

We introduce an aggregate utility function separable in meat consumption, m. The good, m, is itself a constant elasticity of substitution (CES) aggregation of the three meat products considered in this paper: beef (B), pork (P), and chicken (C).

$$u(c, m) = \frac{c^{1-\eta}}{1-\eta} + \alpha_m \frac{m^{1-\xi}}{1-\xi} \tag{6}$$

$$m = \left(\sum_{j \in \{B,P,C\}} \omega_j q_j^{\varepsilon} \right)^{\frac{1}{\varepsilon}} \tag{7}$$

In (6), the parameter ξ measures the diminishing returns to animal products, and α_m scales the total utility from these products. In (7), the parameter ε maps to the elasticity of substitution between these goods, with higher values implying the goods are less substitutable. The ω_j parameter is a taste shifter within the CES, which allows for consumption between these goods to vary even if their prices were equal. The full parameterization, and our methods for deriving them, are detailed in the online appendix (Section C). This function is specified with features—namely, the large marginal utility of m as $m \to 0$—that conservatively anchor the resulting optimal policy to current levels of production. If large reductions are recommended under this specification, there will be little disagreement between this set up and alternative models over the desirability of more modest policy proposals.

With these additions, the framework permits an optimization exercise over the size of animal agriculture. In the initial exercise, we give the planner a simple optimization over uniform reductions across all three products. The optimization problem is static over contemporaneous production. Because utility is separable across time and inputs, and the problem is dominated by animal welfare as opposed to the climate costs, there should not be meaningful differences in the fully dynamic optimization. Figure 3 presents the results along a range of animal utilities.

Optimal reduction, if done bluntly across all animal products, is just under half (45%) in our baseline calibration: the current size of animal agriculture is nearly twice as large as it would be under this formulation of optimal policy that accounts for animal welfare and climate costs. In the functional form chosen, the marginal human-utility of meat eating increases dramatically as its production is restricted—to ∞ as $m \to 0$, in fact. Further, because global meat consumption has increased at

[7] Even without the existence of sizeable animal welfare externalities, global food commodity markets, including those for livestock, are distorted by a complex mix of taxes, subsidies, quotas, and tariffs (OECD 2019). We leave for future research the study of optimal policy design in the face of these other substantial distortions and externalities.

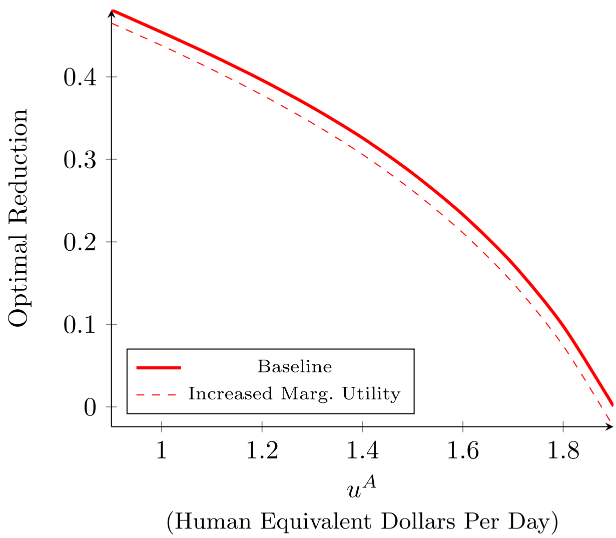

Fig. 3 Optimal reductions in animal agriculture by animal utilities. Optimal animal agriculture as a function of animal utility, u^a. The x-axis reports the human-equivalent-consumption that would generate the u^A used in that scenario, i.e, a value of \$1.50 here implies u^A is equivalent to the utility humans receive living at \$1.50 per day. The "Increased Marg. Utility" curve increases α_m by 10%, demonstrating that this avenue (i.e., a greater taste for meat) has small effects on the optimal policy

a slower rate than incomes in the data, our calibration puts more curvature on utility from meat (m) than consumption (c). This curvature implies that the increase in marginal utility of m happens quickly, restricting the amount the planner is willing to take from humans.

In light of the many layers of uncertainty underpinning this calculation—and the paucity of data on large, voluntary reductions in aggregate meat consumption—we are uncertain about the exact magnitude of optimal reductions, but feel confident that if animals have lives not worth living such reductions would be significant.

Figure 3 takes steps towards demonstrating this by explicitly solving for optimal reductions for different levels of assumed animal welfare. For animal lives that are equivalent to human utility at \$1.40 per day or less, the reductions remain over 30%. As lives become nearly neutral (\$1.80 per day human equivalent utility), optimal reductions are still 10%—a non-trivial number given the currently large size of the animal agriculture sector and its projected growth. Interestingly, when animal lives are neutral, so that the only externality is the environmental impact, optimal reductions are indistinguishable from zero on this plot. Coupled with our conservative modeling choices, this blunt "vegetarian" tool forces reductions of pork and chicken along with beef, making it a poorly targeted climate policy option. In conjunction with Fig. 2, if this axis were extended into the range where animals instead have lives worth living, the planner would recommend large increases in this sector at the expense of a warming planet. For previously stated reasons we are less confident in the reliability of our model in that region of the parameter space, so we note this qualitative difference without an accompanying quantitative exercise.

In addition to robustness along the animal welfare dimension, these optimal reductions remain largely unchanged if consumers' preferences for meat become stronger. When the marginal utility of meat consumption is increased by scaling α_m up by 10% (the dashed line in Fig. 3), optimal reductions are uniformly lowered by approximately two percentage points across the distribution of animal welfare.[8] This is unsurprising given the magnitude of the social costs relative to the plausible private benefits of meat consumption.

4.2 Cross-product results and robustness

Perhaps more decision relevant than the large social costs at the dietary level are the differences in these costs *across* products. These products appear relatively substitutable to consumers (e.g. Schlenker and Villas-Boas 2009), and the large differences in baseline social costs suggest possible welfare enhancing directions of such substitution (namely, from poultry to beef). In the framework of Sect. 4.1, a product-specific optimization would reflect this: if poultry imposes more social costs than beef, the planner will prioritize reductions in the poultry sector (see Appendix A, Table A1). This qualitative result directly relies on cross-product marginal social cost differences, which themselves rely on modeling assumptions regarding animal welfare and climate costs. More informative, then, than formalizing the planner's predictable cross-product prioritization under our baseline assumptions, is to study the model modifications that retain or reject the finding that poultry has the highest social costs.

The approach we take is to map the parameter space into regions where poultry remains more socially costly than beef and regions where the opposite is true. In other words, we document what must be true within our framework to rationalize the prioritization of beef reductions. The focus is on these products because they represent the highest welfare substitution in the baseline and saliently highlight the tension between animal welfare and climate considerations. Generating a reversal of the baseline results—that beef is more socially costly—requires decreasing the size of the animal welfare externality and/or increasing the climate costs of GHG emissions. We proceed by simultaneously adjusting along both dimensions.

The results of this exercise are summarized in Fig. 4. The size of the animal welfare externality, relative to baseline, is plotted on the y-axis. Recall that θ is the welfare weight on animals, which we set at 1.0 in the baseline. Because this parameter enters Eq. (2) multiplicatively, reducing its value to 0.01 is analogous to scaling the animal welfare externality to 1% of its original level. Whether that arises in practice from a difference in this welfare weight or the fact that animal lives may be much closer to neutrality than we assume makes no difference here. Also note that we bound this exercise from below at 0; we continue to doubt the usefulness of the

[8] Note that under the increased marginal utility setting, animal agriculture optimally *increases* when animal lives are neutral. This is because we leave prices unchanged, so the market is temporarily out of (private sector) equilibrium at current production levels.

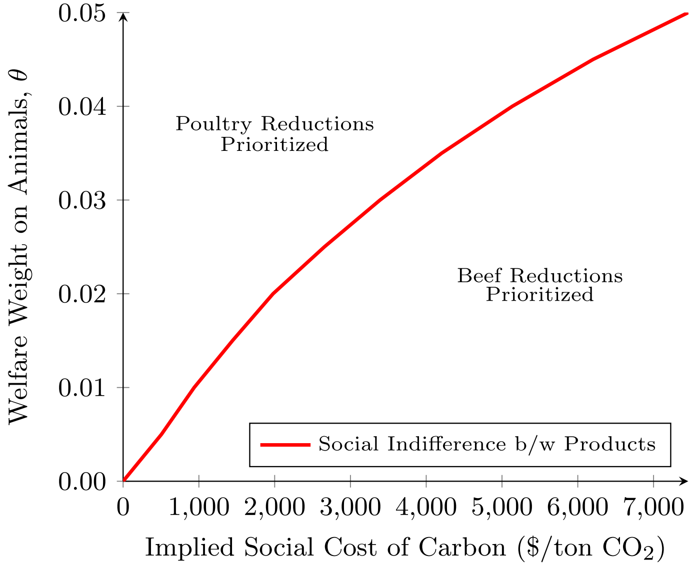

Fig. 4 Product reduction priorities by welfare weight and social cost of carbon. Mapping between optimal product reductions (chicken and beef), the welfare weight on animal utility, and the social cost of carbon (SCC). Areas above the social indifference curve indicate (SCC, θ) combinations such that reducing the global poultry sector is prioritized and conversely for beef below the curve

exercise generally for the case in which the baseline sign on the animal welfare externality is reversed.

The climate dimension is less straightforward. To increase the climate costs, we adjust two dimensions of the model. First, we set the social rate of time preference (ρ) to near-zero (0.001% annually). This reflects a near-equal concern for future generations and is a straightforward and ethically appealing method for increasing the value of climate damages. We then introduce a new channel whereby climate change physically impacts the rate of economic growth—in addition to the standard level effects—à la Moore and Diaz (2015). This too is a straightforward method for increasing climate damages because growth effects compound into the future (now valued at near-present levels with the discount adjustment). The social indifference curve between beef and poultry is found by solving for combinations of animal welfare weights and growth rate damages that equalize the marginal social costs of these products; the regions are then separated by whether they lie above or below this indifference curve. The resulting climate model—with the new discount rate and additional economic damages—is summarized on the *x*-axis by performing one additional step of computing the implied social cost of carbon (per ton of $C0_2$) of this parameter combination. For example, a value of 1000 corresponds to an underlying climate module with an SCC of \$1000 per ton of CO_2.

Figure 4 indicates that the parameter combinations necessary to prioritize beef reductions are extreme relative to the baseline calibration. Even assuming the true animal welfare costs are only 1% of our baseline value, the climate modifications necessary for indifference between beef and poultry imply an SCC of around \$1000 per ton of CO_2. Oft-cited DICE and EPA estimates of the SCC are in the \$30–60 per-ton range (Nordhaus 2017); common arguments for certain damage specifications and discount factors normally increase these estimates into the \$100–\$500 range. We do not endorse a particular SCC. If ethicists are correct that the social rate of time preference ought to be near-zero, a \$1000 per-ton SCC may indeed be possible if damage functions are on the higher end of existing estimates. However, for the social cost of beef to be comparable to that of poultry, it must *also* be the case that we have overestimated animal welfare costs by two orders of magnitude. If we have only overestimated the animal welfare costs by 20-fold ($\theta = 0.05$), the SCC would need to be near \$7000. Given this value is so large relative to any estimate we have seen in the environmental economics literature, we do not attempt to solve this indifference curve for values of $\theta > 0.05$. Despite the restricted domain of this plot, one can see that extending the vertical axis to our baseline value of $\theta = 1$ would illustrate the (very) small parameter space over which beef reductions would be prioritized.

There remain uncertainties not explored in this exercise that can overturn this result. One important source is relative welfare across animals, which remains especially uncertain. If beef cattle had much worse subjective experiences (per unit of time) than chickens, this would be an important consideration pushing against the result that poultry reductions should be prioritized. However, insofar as the literature has addressed animal welfare (e.g. Garnett et al. 2013), most believe beef cattle have better lives than industrial chickens, so we are not especially troubled by this lingering uncertainty. We do not doubt there may be concerns outside of this, but Fig. 4 demonstrates that along the most important dimensions of our model, major modifications are necessary to reverse the cross-product implications of Table 2.

5 Summary and conclusions

Despite recurring discussions and policy debate about farmed animal welfare and GHG emissions from livestock production, there are notably few economic assessments of these costs. We fill this gap by providing a rigorous study of these externalities in a unified setting and find their sum to be very large, driven by animal welfare considerations. Consequently, the current size of animal agriculture, especially the poultry sector, is much larger than it would be under the choice of a benevolent planner with an inclusive welfare function. While these costs rely on an important, untestable, assumption—that a farmed animal's life is subjectively worse than nonexistence—the qualitative result is robust to a wide range of variability in how this assumption is implemented. It is not until farmed animal lives become net-pleasurable that our framework finds that this activity is not a substantial burden on social welfare.

An important limitation to this study, related to the concern about whether the representative animal has a life worth living, regards the differences in welfare across animals, even within species. To the extent that our framework estimates the social costs of a life not worth living, the results are not externally valid for conditions wherein animals do have worthwhile lives. Just because *some* animals have net-negative lives, if even that is true, does not mean our framework recommends reducing the size of operations producing net-pleasurable lives. In terms of practical policy making, the results then only apply to specific industries or firms producing animals with net-displeasurable lives. Design of such policy is beyond the scope of this paper, though we note that a majority of animals raised for food in the developed world reside in industrial farms, and these are the operations many believe do not produce worthwhile lives. This serves as a natural place to turn.

Our results echo the calls made by Johansson-Stenman (2018) and Carlier and Treich (2020) of furthering academic research in this area. While we have not resolved challenging questions at the intersection of ethics and the biological sciences regarding the subjective value of animal lives, we have uncovered what we view as important conditional results: using standard economic techniques, if animal lives are not worth living, there are tremendous social costs generated in this market. The tools of economics do not require precise values on all parameters to produce useful insights.

Accordingly, some specific future research topics that we believe economists could shed light on, if only imprecisely, are optimal investments designed to improve animal rearing conditions (from the point of view of the animal) and optimal consumption under more humane production methods. Moreover, outside of economics, the advancing work on differences in welfare across species (e.g. Schukraft 2020) and production methods should eventually allow researchers to produce more specific policy recommendations. At the very least, our current results highlight the potential first-order nature of animal welfare questions relative to other issues competing for the attention of the welfare economics community.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s00355-023-01451-9.

Acknowledgements This paper was previously circulated as "Optimal Animal Agriculture Under Climate and Population Externalities." We would like to thank Lewis Bollard, Nathan Chan, Matthew Clancy, Zachary Freitas-Groff, Daniel Hellerstein, Nicolas Treich, Le Wang, participants at the Population, Social Welfare, and Climate Change conference at the University of Texas at Austin; the virtual Welfare Economics and Economic Policy seminar series; and the 8th Workshop on Global Priorities Research at the University of Oxford for comments and suggestions that have greatly improved the paper. All remaining errors are our own.

Funding Not applicable.

Data availability All data is available at https://github.com/kevinkuruc/DICEFARM.

Code availability All code is available at https://github.com/kevinkuruc/DICEFARM.

Declarations

Conflict of interest The authors declare no competing interests.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

References

Adler MD (2008) Future generations: a prioritarian view. George Wash Legal Rev 77:1478
Arrhenius G (2000) An impossibility theorem for welfarist axiologies. Econ Philos 16(2):247–266
Bentham J (1789) An introduction to the principles of morals and legislation. T. Payne and Son, London
Blackorby C, Donaldson D (1992) Pigs and guinea pigs: a note on the ethics of animal exploitation. Econ J 102(415):1345–1369
Blackorby C, Bossert W, Donaldson D (1995) Intertemporal population ethics: critical-level utilitarian principles. Econometrica 63(6):1303–1320
Blackorby C, Bossert W, Donaldson DJ (2005) Population issues in social choice theory, welfare economics, and ethics. Cambridge University Press, Cambridge
Budolfson M, Spears D (2019) Optimal climate policy including animal welfare. CFI working paper, Princeton University
Burke M, Hsiang SM, Miguel E (2015) Global non-linear effect of temperature on economic production. Nature 527(7577):235–239
Carlier A, Treich N (2020) Directly valuing animal welfare in (environmental) economics. Int Rev Environ Resour Econ 14:1–40
Dasgupta P (2008) Discounting climate change. J Risk Uncertain 37(2–3):141–169
Diaz RJ, Rosenberg R (2008) Spreading dead zones and consequences for marine ecosystems. Science 321:926–929
Errickson F, Kuruc K, McFadden J (2021) Animal-based foods have high social and climate costs. Nat Food 2(4):274–281
Espinosa R, Treich N (2021) Animal welfare: antispeciesism, veganism and a "life worth living''. Soc Choice Welf 56:531–548
FAO (2021a) Global livestock environmental assessment model. http://www.fao.org/gleam/en/. Accessed 1 Mar 2021
FAO (2021b) Livestock primary: producing animals/slaughtered dataset. http://www.fao.org/faostat/en/#data/QL. Accessed 1 Mar 2021
Fleurbaey M, Van der Linden M (2021) Fair social ordering, egalitarianism, and animal welfare. Am Econ J Microecon 13(4):466–91
Foley JA, DeFries R, Asner GP et al (2005) Global consequences of land use. Science 309(5734):570–574
Garnett T, Appleby MC, Balmford A et al (2013) Sustainable intensification in agriculture: premises and policies. Science 341(6141):33–34
Gerber PJ, Steinfeld H, Henderson B, Mottet A, Opio C, Dijkman J, Falcucci A, Tempio G (2013) Tackling climate change through livestock—a global assessment of emissions and mitigation opportunities. Technical report, Food and Agriculture Organization of the United Nations
Greaves H (2017) Population axiology. Philos Compass 12(11):e12442
Groff Z, Ng YK (2019) Does suffering dominate enjoyment in the animal kingdom? An update to welfare biology. Biol Philos 34(4):40
Howard PH, Sterner T (2017) Few and not so far between: a meta-analysis of climate damage estimates. Environ Resource Econ 68(1):197–225
Hsiung W, Sunstein CR (2006) Climate change and animals. Univ Pa Law Rev 155:1695–1740
Innes GK, Randad PR, Korinek A et al (2020) External societal costs of antimicrobial resistance in humans attributable to antimicrobial use in livestock. Annu Rev Public Health 41:141–157
Johansson-Stenman O (2018) Animal welfare and social decisions: Is it time to take Bentham seriously? Ecol Econ 145:90–103
Kagan S (2019) How to count animals, more or less. Oxford University Press, Oxford
Korsgaard C (2018) Fellow creatures. Oxford University Press, Oxford
Krutilla JV (1967) Conservation reconsidered. Am Econ Rev 57(4):777–786
Lawson N, Spears D (2020) Population externalities and optimal social policy. Working paper
LoPalo M, Kuruc K, Spears D et al (2019) Quantifying India's climate vulnerability. India Policy Forum 15
Matheny G (2003) Least harm: a defense of vegetarianism from Steven Davis's omnivorous proposal. J Agric Environ Ethics 16(5):505–511

Méjean A, Pottier A, Zuber S et al (2020) Catastrophic climate change, population ethics and intergenerational equity. Clim Change 163:873–890
Millar RJ, Nicholls ZR, Friedlingstein P et al (2017) A modified impulse-response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emissions. Atmos Chem Phys 17(11):7213–7228
Moore FC, Diaz DB (2015) Temperature impacts on economic growth warrant stringent mitigation policy. Nat Clim Change 5(2):127–131
National Academies (2017) Valuing climate damages: updating estimation of the social cost of carbon dioxide. Report, National Academies of Sciences, Engineering, and Medicine, Washington, DC
Newbold T, Hudson LN, Hill SL et al (2015) Global effects of land use on local terrestrial biodiversity. Nature 520(7545):45–50
Ng YK (1995) Towards welfare biology: evolutionary economics of animal consciousness and suffering. Biol Philos 10(3):255–285
Nordhaus WD (1992) An optimal transition path for controlling greenhouse gases. Science 258(5086):1315–1319
Nordhaus WD (2017) Revisiting the social cost of carbon. Proc Natl Acad Sci 114(7):1518–1523
Norwood F, Lusk J (2011) Compassion, By the pound: the economics of animal welfare. Oxford University Press, Oxford
OECD (2019) Agricultural policy monitoring and evaluation 2019. Report, Organisation for Economic Co-Operation and Development, Paris, France
Parfit D (1984) Reasons and persons. Oxford University Press, Oxford
Pearce D (2021) Individual and social welfare: a Bayesian perspective. Working paper, New York University
Ripple WJ, Smith P, Haberl H et al (2014) Ruminants, climate change and climate policy. Nat Clim Change 4:2–5
Schlenker W, Villas-Boas SB (2009) Consumer and market responses to mad cow disease. Am J Agric Econ 91(4):1140–1152
Schukraft J (2020) Differences in the intensity of valenced experiences across species. Report, Rethink Priorities
Shields L (2012) The prospects for sufficientarianism. Utilitas 24(1):101–117
Singer P (1975) Animal liberation. Random House, New York
Springmann M, Godfray HCJ, Rayner M et al (2016) Analysis and valuation of the health and climate change cobenefits of dietary change. Proc Natl Acad Sci 113(15):4146–4151
Tännsjö T (2016) It's getting better all the time. In: Barnhill A, Budolfson M, Doggett T (eds) Food, ethics, and society. Oxford University Press, Oxford, pp 362–365
Thompson PB (2020) Philosophical ethics and the improvement of farmed animal lives. Anim Front 10(1):21–28
Wada Y, Van Beek LP, Van Kempen CM et al (2010) Global depletion of groundwater resources. Geophys Res Lett 37(20):L20402
Weitzman ML (2012) Ghg targets as insurance against catastrophic climate damages. J Public Econ Theory 14(2):221–244
Wirsenius S, Hedenus F, Mohlin K (2011) Greenhouse gas taxes on animal food products: rationale, tax scheme and climate mitigation effects. Clim Change 108(1–2):159–184
Zuber S, Asheim GB (2012) Justifying social discounting: the rank-discounted utilitarian approach. J Econ Theory 147(4):1572–1601