OPEN ACCESS

Gain-of-function research and model organisms in biology

Nicholas G Evans [1], Charles H Pence [2]

[1]Department of Philosophy, University of Massachussetts Lowell, Lowell, MA, USA
[2]Institut supérieur de philosophie, Universite catholique de Louvain, Louvain-la-Neuve, Belgium

Correspondence to
Dr Nicholas G Evans, Department of Philosophy, University of Massachusetts Lowell, Lowell, MA, USA; Nicholas_Evans@uml.edu

NGE and CHP contributed equally.

Received 15 December 2022
Accepted 4 May 2023
Published Online First 15 May 2023

To cite: Evans NG, Pence CH. *J Med Ethics* 2024;**50**:201–206.

ABSTRACT

So-called 'gain-of-function' (GOF) research is virological research that results in a virus substantially more virulent or transmissible than its wild antecedent. GOF research has been subject to ethical analysis in the past, but the methods of GOF research have to date been underexamined by philosophers in these analyses. Here, we examine the typical animal used in influenza GOF experiments, the ferret, and show how despite its longstanding use, it does not easily satisfy the desirable criteria for an *animal model*. We then discuss the limitations of the ferret model, and how those epistemic limitations bear on ethical and policy questions around the risks and benefits of GOF research. We conclude with a reflection on how philosophy of science can contribute to ethical and policy debates around the risks, benefits and relative priority of life sciences research.

INTRODUCTION: A POLICY STORM IN A RESEARCH TEACUP

In 2011, two experiments on influenza were announced for publication in *Nature* and *Science*, describing work by teams at the University of Wisconsin-Madison in the USA and Erasmus University in the Netherlands, respectively. Both studies described the generation of novel strains of influenza—one, a recombinant version of HA H5N1 'avian influenza', the other a reassortant strain of influenza created from both H5N1 and the 2009 H1N1 pandemic 'swine influenza' strain—that were transmissible between mammals.[1 2] The generation of H5N1 influenza viruses that transmitted between mammals had never been accomplished, in laboratory or in nature. And while seasonal influenza has a case fatality rate of approximately 0.1%, H5N1 has killed approximately 60% of the people who have been contracted and is estimated to have a case fatality rate of 14%–33% if it ever caused a pandemic.[3]

Controversy emerged because experts and policymakers were divided as to the ethical justifications for conducting or publishing these 'gain-of-function' (GOF) studies. Security-minded folk pointed out that the two experiments were first and foremost a blueprint for potential bioterrorists who wanted to create a devastating and indiscriminate weapon.[1] Over the coming decade, the debate would expand to influenza and coronavirus research, and focus on the safety issues that would arise as the technique proliferated across the world and repeat experiments threatened a laboratory accident seeding a global pandemic.[4 5]

But a significant proportion of the scientific community claimed in the first place that they would not withhold the research, and second that the research not only had scientific but *social* value. The former claim came down to claims about the value of scientific freedom and the importance of norms of openness in scientific progress.[6] Proponents also claimed that the study of influenza in mammals—and in ferrets in particular—was essential to maintaining surveillance for potential disease pandemics in nature; was useful for developing medical countermeasures against influenza and could aid the development of vaccines. The same proponents, however, would attempt to play down the security and safety fears surrounding GOF experiments by noting that the development of novel strains of influenza that were transmissible between ferrets did not mean those same strains would transmit between humans.[7]

This may strike the reader as scientists attempting to have their cake and eat it too. The strength of the connection between ferrets as an *experimental organism* and humans as its target motivates the experiment, but that connection is not strong enough to motivate fears that the results of the studies are directly transferable to humans. Rather than a mere rhetorical device, however, we think this tension points to deeper questions about influenza studies, and about the role of ferrets within them.

Given the stakes of GOF research, a more careful analysis of the cost–benefit calculus surrounding GOF in influenza is necessary.[5 8] It is further clear that such a cost–benefit analysis needs to carefully consider not only the ethical stakes, but also the *epistemic* stakes. What kinds of scientific knowledge could really be generated by GOF research, and what would the cost of foregoing that knowledge be? As we have just seen, performing that analysis and taking the epistemic stakes seriously require that we explore the system in which that knowledge is generated: the ferret.

The case of the ferret model is interesting for philosophers and bioethicists in at least two complementary ways. First, while model organisms are a live subject of debate in philosophy of science, the literature has not thus far examined ferrets, either in the descriptive or in an evaluative dimension. One might ask, for instance (following a distinction common in this literature[9 10]), whether ferrets are a genuine *model* organism, that is, taken to support a host of broad inferences from ferrets to other mammals, supported by standardised protocols, databases and other epistemic resources, or whether they are instead an *experimental* organism, a special-purpose tool for studying a single phenomenon but lacking that epistemic and infrastructural support. Second, the ambiguities surrounding ferrets provide an opportunity for philosophy of science to

contribute, in concrete terms, to a live ethical and policy issue. Clarifying what ferrets can and cannot do for science—taking the epistemic element of the cost–benefit calculus seriously—allows us to answer questions about the utility of studies that use ferrets.

As a preliminary move we assume that at least in principle, some instances of GOF research are justified. While there is extensive debate on the absolute merits of GOF research that may result in potential pandemic pathogens and the availability of alternatives that might replace GOF as a method in whole or part, we set this debate aside for this work. Rather, we are interested in cases of GOF research in influenza that have a permissible, or at least contestably permissible risk–benefit ratio, and investigate what the epistemic status of the ferret model can tell us about our decisions in these cases.

THE FERRET MODEL OF INFLUENZA

The ferret has long been linked with the study of influenza. After an influenza epidemic in 1933, Smith *et al* at the National Institute for Medical Research's Farm Laboratories in Mill Hill, outside of London, collected throat washings from several patients, and proceeded to assay every available experimental animal model to see if any of them would become infected.[11] By a stroke of luck, a population of experimental ferrets was being maintained to study canine distemper virus, and two ferrets tested began to demonstrate all the classic symptoms of the influenza: fever, food avoidance, weight loss, sneezing ('the ferret has an exquisite sneeze reflex' reported one researcher),[12] fatigue and a runny nose.[13] The work of Smith *et al*, in the end, led to the first-ever isolation of the influenza virus, which had been suspected but not confirmed to be the causative agent of influenza.

What makes the ferret a particularly useful model is it is the only known model which 'can present both the pathogenic and transmissible features of influenza virus infection'.[14] That is, ferrets will both present a course of disease symptoms like humans, giving us an idea of whether or not a particular strain of influenza would lead to severe disease. But they also spread the disease via both droplet and aerosol transmission, giving us a way to determine whether a strain would be highly transmissible among humans. As these are two variables that proponents of GOF claim are necessary to determine the risk of a human respiratory pandemic,[8 15–20] the ferret model has become increasingly important in the study of pandemic influenza.

All model organisms for influenza miss at least one of these other features. Mice require specially adapted influenza viruses and thus cannot be used to directly test pathogenicity of viruses that could infect humans, and do not transmit influenza via aerosols or droplets.[21] One strain of guinea pigs seems to mimic human transmission behaviour,[21 22] but does not exhibit clear signs of infection. Recently, transgenic mice have been developed to identify potential immune evasion of influenza viruses in humans,[23] but these still are used for infection but not transmission studies.

Dietrich *et al*[9] have argued for a collection of 20 different features they argue shape the choice of model organisms in the life sciences. These should not be interpreted as a 'checklist' for a 'good' model organism—rather, they are something more like an overlapping collection of virtues for which scientists typically argue when they defend the use of a model organism. Three of these features are commonly cited in scientific discussions of the ferret model in influenza research:

1. *Phenomenal access* to the relevant features of influenza: 'in the sense of instantiating its typical features or providing insights that can be used towards understanding the phenomenon in question'.
2. *Translational potential*: ferrets have relevant 'physiological or genetic resemblance to humans'[9]—a common sialic acid binding site in the upper respiratory tract, where influenza viruses most often infect humans, where avian hosts have a binding site predominantly found in the human lower respiratory tract.[24]
3. *Responsiveness:* ferrets offer better opportunities for the experimental manipulation of features of interest to researchers,[9] primarily their respiratory anatomy, which exhibits easy access to both upper and lower respiratory features.[25]

That is, ferrets not only give us valuable information about the world, but they are also easy for scientists to study in useful and interesting ways.

But the literature about the use of ferrets as a model for the study of influenza is remarkably self-reflective. In the terms of the features of animal models, scientists in the ferret literature recognise (and lament!) that their systems *fail* to instantiate at least seven of the other virtues that are common in other model organism research (eg, mice, yeast, zebrafish). Ferrets often present difficulties for both *ease of supply* and *financial costs*, as initial costs are higher, standard equipment for the husbandry of ferrets is not widely available, inbred strains have not been developed and influenza is occasionally already endemic in breeders' populations of ferrets.[12 25 26] Animal *ethical considerations* are more significant in ferrets than many other animals used in research—not because they are taken to be a more ethically sensitive subject, but rather due to a lack of 'best practices' that one might find in systems such as mice, slowing animal ethics committee approval.[26]

Standardisation is often lacking, as highly inbred strains of ferrets have not been developed in the way that they have been in rodents, leading to a largely unknown amount of genetic diversity between populations that hamper the kinds inference scientists can make in influenza research.[25] The *viability and durability* of ferrets is often limited, as their expense often means that sample sizes of as few as five animals are regularly used,[26] and stocks are not usually maintained by the laboratories performing the research. Sample sizes in the single digits—particularly when sampling against a population with an unknown amount of genetic diversity in the absence of standardised, inbred strains—add a currently entirely unknown quantity of further uncertainty to research results.

Both the *availability of methods* and *techniques*, as well as our extant body of *epistemic resources* for ferret experiments, are also underdeveloped. On the practical side, viral inoculation methods, study endpoints and necropsy methods are not necessarily shared across all laboratories.[27] Ferret-specific reagents are often commercially unavailable.[13] More broadly, the ferret genome was only published in 2014[28]—notably, after the controversial GOF research discussed above was announced in late 2011.[7]

In short, ferrets are a difficult, expensive and minimally standardised system. The simple existence of scientific worry does not mean that ferrets somehow 'fail' to be a model organism; however, rather, the overall summaries of the effectiveness of the ferret model are often quite nuanced, even those written by virology researchers themselves. Belser *et al* write that 'these confounders result in heterogeneity with regard to procedures and practices established at all levels of research, from individual

investigators or institutions to broad country-specific regulations'.[13 14] Oh and Hurt write:

> The use of ferrets for influenza studies has been limited by factors such as animal availability, genetic heterogeneity (out-bred), the requirement of a complex husbandry facility and caging system, and a lack of immunological reagents and genetically modified mutants for immunological investigation. […] Ideally, a larger number of ferrets should be used but limitations such as high experimental cost, low animal availability, limited caging capacity and ethical constraint, typically restricts most studies to group sizes of five or less ferrets.[26]

Ferrets are therefore, at the very least, a peculiar experimental organism, peculiar enough to justify an exploration of whether the kinds of epistemic limitations that ferret researchers have mentioned should lead us to re-evaluate the role of ferrets in influenza research and the kinds of conclusions that are drawn from them. While remaining the only system capable of modelling both transmission and infection behaviours in influenza research, the scientific community itself recognises that the system has a host of quirks and flaws. Of course, ferrets are not particularly unusual in this regard. Scientists will often have very good reasons to use what might appear from the outside to be 'better' and 'worse' model organisms in different contexts, including—importantly—there being no better model available for an important set of research questions, as in influenza. But these characteristics of the model—both 'positive' and 'negative', 'fundamental' and 'practical'—have epistemic impacts that will condition the nature of the knowledge that we can derive using ferrets and, by extension, what we might hope to learn from GOF research.

FERRETING OUT THE IMPACT OF EPISTEMIC UNCERTAINTY ON GOF

Recognising the limitations and advantages of the model system allows us to begin to chart how this varied epistemic uncertainty in ferrets should influence the cost–benefit analysis surrounding GOF research. The risk and benefit assessment of GOF research commissioned by the National Institutes of Health, for example, in evaluating whether GOF research can answer the question 'can animal influenza viruses become transmissible between humans', simply states that 'the key limitations of this approach are that observations in animal models may not translate to humans and that the adaptive changes observed in the laboratory may not be possible in nature'.[5] While this reflects part of the kinds of uncertainty developed in the last section, such an assessment (common to essentially any animal research) does not do justice to the rich landscape of uncertainty we described there. Likewise, work on the ethics of GOF research has never addressed the methods of GOF as animal research in detail.[29 30]

Even when it is discussed, the epistemic uncertainty surrounding the scientific knowledge developed in GOF contexts is often not developed fully. For instance, writing about the regulation of GOF experiments, Casadevall *et al* claim 'the debate [around GOF] has largely ignored the question of the epistemological value of such experiments'. But when they turn to the nature of the knowledge that GOF is taken to produce, they write that GOF experiments can be defended 'because they yield information that is consistent with the normative standards of the fields of microbiology and infectious diseases, and as such, they provide information that is immediately accessible and interpretable in the context of standards in the field'.[15]

This evaluation is, of course, correct: articles surrounding GOF research adhere to the kinds of best practices that the field has developed as a whole; there is no reason to believe that they have not been peer reviewed in detail; and they are highly cited and widely discussed. But these social and contextual facts do not amount to the kind of evaluation of the 'epistemological value' of experiments in ferrets for which Casadevall *et al* called in the first place. Researchers working within virology itself recognise that there are some good reasons to believe that ferret research might not offer us the kinds of quality guarantees available in other domains in the contemporary life sciences, and the risk and benefit assessment mentioned above noted, for numerous scientific applications of GOF research, that the social and contextual facts 'are constrained by scientific uncertainties associated with the data'.[5]

Evaluating the impact of the scientific challenges in influenza research on the cost–benefit analysis for GOF research is not an easy task. Assessing the quality of individual GOF papers (especially as an outsider in the field) is likely to be somewhere between impossible and counterproductive. We can instead attempt to take a higher-level approach: what kinds of epistemic uncertainty might one expect to find in a research programme with the sorts of acknowledged concerns we canvassed in the last section?

Recall just what it is that GOF research is supposed to provide us with knowledge about. We are attempting, essentially, to predict the evolutionary future because doing so is instrumentally valuable to preventing a catastrophic disease pandemic.[1] What kinds of mutations might plausibly arise within a given viral lineage, and how might those mutations lead to changes in the viruses' ability to infect and transmit between humans? In several cases, we have successfully made such predictions, testing, for instance, certain kinds of mutations in GOF contexts that were later discovered in human influenza viruses, which went on to be incorporated into the influenza vaccine.[19] Any critique of this research must, to be sure, take this success into account.

Nonetheless, the kinds of concerns that we have already seen with the methodologies adopted in and the statistical power of ferret-based research remain important sources of epistemic uncertainty. We can see this in at least two different ways (in addition to the uncertainty arising from sample sizes already discussed above). First, results arising from a GOF research programme will require an extensive amount of comparison against baselines—we are fundamentally in pursuit of *enhanced* virulence or transmission with respect to some external standard. But our capacity to compare results across laboratories is precisely one of the points of friction most identified by virologists working in ferrets: given the differences between ferret populations and the myriad ways to inoculate ferrets with an influenza virus, to measure viral endpoints and to detect transmission, being certain that such a comparison of viral characteristics is legitimate, are no easy matter.

Finally, in the context of many particular GOF research programmes, this uncertainty in particular results is magnified yet further by the overall experimental design. GOF experiments like those in avian influenza discussed above turn not only on the low-level results in ferrets, but also on a number of other uncertain experimental choices: we must select a strain of avian influenza on which to work (among the many known and likely many more unknown such strains), and we often target a particular site at which to mutate that strain (sometimes drawn from prior experience with other influenza viruses, sometimes from our biochemical knowledge of relevant protein binding). Both of

these are, in essence, gambles on the correct initial conditions for our future evolutionary predictions and magnify the error bars (so to speak) on our final conclusions.

Previous work has maintained that some of these other conditions entail the existence of reliable alternatives to GOF research. Lipsitch and Galvani, in particular, have noted that epistatic gene–gene interactions within influenza viruses provide a good reason to prefer other methodologies to GOF research.[31] That is, identifying the set of genes X in H5N1 highly pathogenic avian influenza as associated with mammalian transmission is not sufficient to show that all H5N1 viruses with X are mammalian transmissible. This is because X interacts with other, sometimes unknown, genes, which may undo its transmission potential. So X indicates mammalian transmissibility except in the case of sets of genes Y, Z, etc.

Evans likewise has argued that sometimes the stated questions GOF is intended to ask do not establish that GOF is uniquely suitable to answer those questions.[4] He argues that if the aim, as in the Fouchier GOF paper, is among other things to find an H5N1 highly pathogenic avian influenza virus whose HA arm (the segment of influenza virus that binds to and enters host cells) is modified sufficiently to bind to receptors in human lungs, much of this work could be answered using attenuated viruses. That is, if the aim is to show that some HA binds to mammalian cells, then a pathogenic influenza virus is not necessary from the outset—if not none, then less GOF research could be done.

However, even with these alternatives, the risk–benefit assessment performed on GOF research identified nine potential experimental aims, particularly around the potential for reassortment of viruses, which would benefit uniquely from GOF methods[5] (while the journal rarely uses page references, we note this is a >1000 page document—the relevant tables are pp. 252–254 of the full report). With this in mind, even if we replaced all GOF research that could be suitably performed using alternate methods, the remainder still are mired in the epistemic uncertainty we identify here. We thus join with an increasing chorus of authors arguing that we should explore more seriously the possibility of non-GOF alternatives that could, at least potentially, obtain us the same kinds of benefits while radically reducing the potential costs.[4 31] But second, and again in line with the epistemic focus of the rest of our work here, we want to consider the potentially damaging *epistemic* effects of high-stakes, high-uncertainty research in cases in which GOF is at least in principle justified.

THE RISK OF BEING WRONG

There are several attendant risks to being wrong about the epistemic benefits of GOF research. The first and most obvious is that if we are wrong, the social benefits of these studies might never be realised. That is, the compelling reason given to expose others to the risks of potential pandemic pathogens might not be obtained. Rather than finding a novel combination of genes in influenza that signifies pandemic potential, we have instead created a novel virus with enhanced features while knowing nothing true about why it is or is not a potential pandemic virus.

One question that arises here is whether we can expect this change in risk to be *symmetric*. That is, if the knowledge gained from GOF ferret studies is limited, there are reasons to believe that in some cases, the reason our knowledge is limited also implies that the risk of harm from GOF research might also be attenuated. If, for example, we cannot extrapolate our knowledge about the role of the HA arm of influenza, which tracks its capacity to infect human cells, from ferrets to humans, then we cannot use this research to generate novel vaccine candidates. But likewise, if the relation between infection in ferrets and in humans is not straightforward, we may also have less reason to believe that this research could, through a laboratory accident for example, lead to the release of a novel disease in humans that seeds a global pandemic. In the event these risk changes are fully symmetric, then, our answers to the epistemic questions would not change the overall valence of our risk–benefit assessment: the expected benefits go down, sure, but so do the risks.

But there are much stronger reasons, we think, to believe this change is sometimes *asymmetric*. That is, the epistemology of GOF research might give us reason to doubt our ability to realise its purported social value, while the harm it poses remains unchanged (or is reduced much less than are its benefits). We might expect this to arise in serial passaging studies such as occurred in the Fouchier paper that generated so much controversy in 2011. Recall that in addition to the uncertainty associated with ferrets, the information gleaned about the genetic basis for mammalian transmission of highly pathogenic avian influenzas is incomplete owing to a lack of comparative knowledge about the epistatic interactions between different genes in the influenza genome itself.[31 32] However, the experiments *did* result in a mammalian-transmissible form of influenza, which in principle could spread in humans. So whether or not ferrets give us the right information we need to know *why* pandemic influenzas arise that are transmissible in humans, and defend against them, we have still constructed a mammalian transmissible strain of influenza with its attendant risks.

There might, however, be a more interesting way in which the risks of being wrong from GOF studies are important. If the problems with ferrets are as deep as is suggested in the literature, then GOF studies—and many other influenza studies besides—are already engaged in a process of muddying the virological waters, merely in virtue of using their preferred animal organism. In the study of epistemic communities, for example, it has been established that patterns emerge when productive and *popular* scientific discoveries engender followers who repeat similar patterns of work as those who came before.[33] In the study of the genomic determinants of human transmissibility of avian influenza viruses, it has been noted by others that GOF research ballooned after 2011 as a popular method, in part because it appeared to have been a success (and presumably also because it was being published in premier outlets). Yet, if ferret-based GOF research is as epistemologically tenuous as ferret researchers themselves sometimes indicate, then there are reasons to believe that errors introduced by GOF research that are derived from the epistemic limitations of ferrets may ultimately lead to more, and not less, confusion. It would behove the scientific community to better establish the model of the ferret so that it can be determined if this direction for the community is indeed valuable in the way supposed, or if resources are being poured into scientific dead ends.

Here, then, the social implications of epistemic uncertainty could be quite serious. Research on the virological characteristics of influenza could produce false positives, or send researchers down the wrong path to knowledge about the virus' properties. This could undermine the public health aims that GOF points at, such as surveillance, or vaccine and medical countermeasure development. Here, the risk of being wrong is not simply the risk of being wrong for a particular study in isolation, but the risk that being wrong alters our epistemic landscape, and creates poor priority setting in the division of our cognitive labour in the life sciences in aid of practical projects.

The case of GOF research thus gives us a concrete opportunity to connect work on biomedical and research ethics, traditional questions in the philosophy of science and the philosophy of biology, and studies of transmission and construction of knowledge under risk and uncertainty—and all these in an environment of potential genuine social benefit. We hope that our efforts here will encourage constructive contributions by philosophers of science to this rapidly evolving, complex and important cluster of epistemological concerns.

PROSPECTS FOR POLICY

More than an interesting epistemological feature of influenza research, uncertainties in ferret models have potentially important implications for bioethically guided policies that govern biosecurity and biosafety.

At the local level, individual researchers might do more, as part of their existing research programmes, to refine and standardise the ferret model of influenza research. Time spent sequencing ferret genomes and developing more standardisation (or at least documentation) of the influence of ferret genetics on influenza research would improve research. The choice to use seropositive or (rarer) seronegative ferrets might be made explicitly—seronegative ferrets allow for naïve infection of influenza, where seropositive ferrets more closely mirror humans (who are infected and reinfected by different variants of the virus).

At the level of the review of individual projects, existing policies that provide oversight of GOF research might include review of the value of the projects in light of ongoing concerns with the experimental animal of choice. Projects that do not include methods to mitigate uncertainty, or acknowledge it as part of experimental design, might be returned or required to substantiate their efforts. Given the potential asymmetry of value in GOF research, this might be a priority for risky research where less risky influenza research can afford the epistemic ambiguities brought on by ferrets.

At the level of funders, more might be spent explicitly on developing a model organism to optimise the use of scientific resources.[34] Here, the social value of the work is clear: as the best, if not only model for influenza transmission research, the improved epistemic status of ferret research entails the improved epistemic status of public health responses to influenza. This would benefit GOF research, yes, but also all influenza research that uses ferrets as a model for humans.

One potential consequence of this last option that might occur to the reader (and suggested to us by a reviewer) first is our analysis might entail more, and not less, GOF research. If indeed there are methodological gaps in ferret research, then this would naturally entail more research to better establish the ferret as a model, some of which might be (or require) more GOF research.

But this connection is not as straightforward as it appears. More research on the ferret model would be valuable. And with this, we could plausibly see a call for increase in virology funding. But it is unlikely to be GOF research. The problems of ferret research are much broader than just GOF research, and are thus not necessarily solvable by GOF research itself. Better improving the physical and epistemic infrastructure of the ferret such as systematising its genetics, better verifying its prediction of transmissibility and developing better accounts of sampling for different kinds of virology research would be immensely useful to virology, and arguably to influenza research with strong public health aims. But none of this need be GOF, and it may in fact be a misuse of resources to perform GOF attempting to aim at better ferret research.

Second, some of the questions about ferrets as models are scientifically complex and would require questions of professional standards that precede GOF entirely. Take the above example of seronegative versus seropositive ferrets. This is a question for the professional community that may need collective action but not research per se. But if research is needed, then it would undoubtedly include examining *which influenza strain ferrets should come with* in order to best advance the epistemic and public health aims of influenza research, a question that does not easily or obviously lend itself to GOF research.

Whether at the level of individual researchers, professional practice or a new funding programme to better understand the ferret, the goal our conclusion entails is a common one. That is, if we take seriously the importance of the ferret to pandemic preparedness, then we should ultimately elevate it to model organism status in the same way that mice are to cancer research (among other fields). This would be a considerable scientific achievement, and the development of a necessary piece of scientific infrastructure to respond to a future influenza pandemic. And it would better position GOF research to be taken seriously as a, all other considerations being equal, part of that preparedness effort.

Twitter Nicholas G Evans @neva9257

Contributors NGE and CHP both conceived of the project. CHP wrote the first draft of the first half of the paper, NGE, the second half. Both authors revised the draft into a coherent whole, with NGE providing final editing and submission. Both authors contributed equally to discussing reviewer comments and revising the manuscript accordingly. NGE acts as guarantor for this research.

Funding This study was funded by Greenwall Foundation (no award/grant number), Davis Educational Foundation (no award/grant number) and Air Force Office of Scientific Research (FA9550-21-1-0142).

Competing interests None declared.

Patient consent for publication Not required.

Ethics approval Not applicable.

Provenance and peer review Not commissioned; externally peer reviewed.

Data availability statement No data are available.

ORCID iDs
Nicholas G Evans http://orcid.org/0000-0002-3330-0224
Charles H Pence http://orcid.org/0000-0002-6836-6047

REFERENCES

1 Evans NG. Great expectations -- ethics, avian flu and the value of progress. *J Med Ethics* 2013;39:209–13.
2 Evans NG, Lipsitch M, Levinson M. The ethics of biosafety considerations in gain-of-function research resulting in the creation of potential pandemic pathogens. *J Med Ethics* 2015;41:901–8.
3 Li FCK, Choi BCK, Sly T, *et al*. Finding the real case-fatality rate of H5N1 avian influenza. *J Epidemiol Community Health* 2008;62:555–9.
4 Evans NG. Ethical and philosophical considerations for gain-of-function policy: the importance of alternate experiments. *Front Bioeng Biotechnol* 2018;6:11.
5 Gryphon Scientific. Risk and benefit analysis of gain of function research. Washington, DC Grphyon Scientific; 2016. Available: http://osp.od.nih.gov/sites/default/files/Risk%20and%20Benefit%20Analysis%20of%20Gain%20of%20Function%20Research%20-%20Draft%20Final%20Report.pdf
6 Evans NG, Selgelid MJ, Simpson RM. Reconciling regulation with scientific autonomy in dual-use research. *J Med Philos* 2022;47:72–94.
7 Enserink M. Scientists brace for media storm around controversial flu studies [Sci. Insid]. 2011. Available: https://www.science.org/content/article/scientists-brace-media-storm-around-controversial-flu-studies
8 Casadevall A, Imperiale MJ. Risks and benefits of gain-of-function experiments with pathogens of pandemic potential, such as influenza virus: a call for a science-based discussion. *MBio* 2014;5.

9 Dietrich MR, Ankeny RA, Crowe N, *et al*. How to choose your research organism. *Stud Hist Philos Biol Biomed Sci* 2020;80:101227.
10 Ankeny RA, Leonelli S. Model organisms. *Model Organisms*. Cambridge: Cambridge University Press, 2020.
11 Smith W, Andrewes CH, Laidlaw PP. A virus obtained from influenza patients. *Lancet* 1933;222:66–8.
12 Radigan KA, Misharin AV, Chi M, *et al*. Modeling human influenza infection in the laboratory. *Infect Drug Resist* 2015;8:311–20.
13 Belser JA, Eckert AM, Tumpey TM, *et al*. Complexities in ferret influenza virus pathogenesis and transmission models. *Microbiol Mol Biol Rev* 2016;80:733–44.
14 Belser JA, Katz JM, Tumpey TM. The ferret as a model organism to study influenza A virus infection. *Dis Model Mech* 2011;4:575–9.
15 Casadevall A, Howard D, Imperiale MJ. An epistemological perspective on the value of gain-of-function experiments involving pathogens with pandemic potential. *MBio* 2014;5:e01875-14.
16 Imperiale MJ, Casadevall A. A new approach to evaluating the risk-benefit equation for dual-use and gain-of-function research of concern. *Front Bioeng Biotechnol* 2018;6:21.
17 Imperiale MJ, Howard D, Casadevall A. The silver lining in gain-of-function experiments with pathogens of pandemic potential. In: Yamauchi Y, ed. *Influenza Virus: Methods and Protocols*. New York, NY: Springer, 2018: 575–87.
18 Fouchier RAM, García-Sastre A, Kawaoka Y, *et al*. Transmission studies resume for avian flu. *Science* 2013;339:520–1.
19 Schultz-Cherry S, Webby RJ, Webster RG, *et al*. Influenza gain-of-function experiments: their role in vaccine virus recommendation and pandemic preparedness. *MBio* 2014;5:e02430–14.
20 Adam DC, Magee D, Bui CM, *et al*. Does influenza pandemic preparedness and mitigation require gain-of-function research? *Influenza Other Respir Viruses* 2017;11:306–10.
21 Bouvier NM. Animal models for influenza virus transmission studies: a historical perspective. *Curr Opin Virol* 2015;13:101–8.
22 Lowen AC, Mubareka S, Tumpey TM, *et al*. The guinea pig as a transmission model for human influenza viruses. *Proc Natl Acad Sci U S A* 2006;103:9988–92.
23 Nguyen T-Q, Rollon R, Choi Y-K. Animal models for influenza research: strengths and weaknesses. *Viruses* 2021;13:1011.
24 Richard M, van den Brand JMA, Bestebroer TM, *et al*. Influenza A viruses are transmitted via the air from the nasal respiratory epithelium of ferrets. *Nat Commun* 2020;11:766.
25 Maher JA, DeStefano J. The ferret: an animal model to study influenza virus. *Lab Anim* 2004;33:50–3.
26 Oh DY, Hurt AC. Using the ferret as an animal model for investigating influenza antiviral effectiveness. *Front Microbiol* 2016;7:80.
27 Belser JA, Eckert AM, Huynh T, *et al*. A guide for the use of the ferret model for influenza virus infection. *Am J Pathol* 2020;190:11–24.
28 Peng X, Alföldi J, Gori K, *et al*. The draft genome sequence of the ferret (Mustela putorius furo) facilitates study of human respiratory disease. *Nat Biotechnol* 2014;32:1250–5.
29 Resnik DB. H5N1 avian flu research and the ethics of knowledge. *Hastings Cent Rep* 2013;43:22–33.
30 Douglas T. The dual-use problem, scientific isolationism and the division of moral labour. *Monash Bioeth Rev* 2014;32:86–105.
31 Lipsitch M, Galvani AP. Ethical alternatives to experiments with novel potential pandemic pathogens. *PLOS Med* 2014;11:e1001646.
32 Yamauchi Y. Influenza virus. In: Yamauchi Y, ed. *Why Do Exceptionally Dangerous Gain-of-Function Experiments in Influenza*. New York, NY: Springer, 2018: 589–608.
33 Muldoon R. Diversity and the division of cognitive labor. *Philos Compass* 2013;8:117–25.
34 Buchanan A, Kelley MC. Biodefence and the production of knowledge: rethinking the problem. *J Med Ethics* 2013;39:195–204.