

# Eliciting judgements about dependent quantities of interest: The SHELF extension and copula methods illustrated using an asthma case study

Björn Holzhauer<sup>1</sup>    Lisa V. Hampson<sup>1</sup>    John Paul Gosling<sup>2</sup>  
Björn Bornkamp<sup>1</sup>    Joseph Kahn<sup>3</sup>    Markus R. Lange<sup>1</sup>  
Wen-Lin Luo<sup>3</sup>    Caterina Brindicci<sup>1</sup>    David Lawrence<sup>1</sup>  
Steffen Ballerstedt<sup>1</sup>    Anthony O’Hagan<sup>4</sup>

<sup>1</sup>Novartis Pharma AG, Basel, Switzerland

<sup>2</sup>JBA Risk Management Ltd, Skipton, United Kingdom

<sup>3</sup>Novartis Pharmaceuticals Corporation, East Hanover, NJ, USA

<sup>4</sup>The University of Sheffield, School of Mathematics and Statistics,  
Sheffield, United Kingdom

## Abstract

Pharmaceutical companies regularly need to make decisions about drug development programs based on the limited knowledge from early stage clinical trials. In this situation, eliciting the judgements of experts is an attractive approach for synthesising evidence on the unknown quantities of interest. When calculating the probability of success for a drug development program, multiple quantities of interest — such as the effect of a drug on different endpoints — should not be treated as unrelated.

We discuss two approaches for establishing a multivariate distribution for several related quantities within the SHEffield ELicitation Framework (SHELF). The first approach elicits experts’ judgements about a quantity of interest conditional on knowledge about another one. For the second approach, we first elicit marginal distributions for each quantity of interest. Then, for each pair of quantities, we elicit the concordance probability that both lie on the same side of their respective elicited medians. This allows us to specify a copula to obtain the joint distribution of the quantities of interest.

We show how these approaches were used in an elicitation workshop that was performed to assess the probability of success of the registrational program of an asthma drug. The judgements of the experts, which were obtained prior to completion of the pivotal studies, were well aligned with the final trial results.

# 1 Introduction

The decision to continue or stop the development of a new drug is an example of high-stakes decision making in the pharmaceutical industry. To continue usually means a commitment to large and costly clinical trials that may expose the enrolled patients to risks, while to stop may mean a missed opportunity to help patients. At the same time, only limited data are usually available. Thus, improving the decision making in these situations is an important problem.

For decision making with no or limited directly relevant data, eliciting the judgements of a group of experts is one approach to effectively combining the available direct and indirect evidence. Expert knowledge elicitation is the process of capturing expert knowledge about one or more uncertain quantities in the form of a probability distribution. It is an important tool to provide understanding of uncertain phenomena and inputs to decision-making processes. There has been a steadily growing demand for elicitation in many fields throughout industry, government and science — see, for example, Garthwaite et al.<sup>1</sup>, Gosling et al.<sup>2</sup>, Usher and Strachan<sup>3</sup>, and Bamber et al.<sup>4</sup>. In particular, elicitation has been advocated and used in pharmaceutical science<sup>5;6</sup> and public health<sup>7;8</sup>. Due to the cognitive biases that experts are subject to, several frameworks and procedures have been proposed to guide the elicitation process in order to minimise these biases. Three leading elicitation protocols are set out and contrasted in the European Food Safety Authority guidance on expert knowledge elicitation in food and feed safety risk assessment<sup>9</sup>. Our preference is for the SHEffield ELicitation Framework (SHELF) described in Section 3 of this paper. The most important reason for this choice is that we believe that SHELF’s unique form of facilitated discussion between the experts provides considerable added value to the elicitation. Another reason, which is important in pharmaceutical applications such as the asthma study described in Section 4, lies in the flexible tools for eliciting dependence between outcomes provided in the latest release of SHELF.

Dependence is an issue when the likely values of a quantity of interest (QoI) depend on another uncertain quantity, because it can be challenging to elicit judgements about the QoI conditional on the other quantity. Similarly, QoIs are often likely to be dependent, in which case the challenge of eliciting a joint distribution for several QoIs arises. There are many methods in the literature for capturing knowledge about dependencies between multiple variables<sup>10;11</sup>. However, these methodologies are typically reported in the literature as standalone methods rather than forming part of a complete elicitation protocol like SHELF. Also, whereas SHELF is a generic protocol that is applicable to a very wide range of applications, most of these methodologies have considerable restrictions.

- They may constrain the type of variables and distributions to be fitted — for example, to Dirichlet distributions for proportions<sup>12;13</sup>.
- They may be tailored for a specific application — for example, land cover<sup>14</sup> or system reliability<sup>15</sup>.

- They may consider complex restructuring for large numbers of dependent variables<sup>16;17;18</sup>.

We present generic methods for eliciting joint distributions through judgements that experts can realistically make. Like the SHELF protocol itself, these methods are applicable in all areas where elicitation is required, and to use them effectively there are important choices to be made. Examples of its use, and the choices made, in any specific field can therefore serve as valuable guides for others to follow. We illustrate their use within SHELF in a pharmaceutical example that we fully describe in Section 2. The example concerns assessing the probability of success (PoS) of a Phase 3 drug development program. Such programs are expensive, resource-intensive long-term commitments for any organisation. The decision to proceed with a Phase 3 program depends on many considerations including the unmet medical need and market opportunity a new drug may address, as well as the probability of success to address these needs. As part of a pilot project to evaluate a new PoS framework at Novartis, we conducted a PoS assessment for an asthma drug. While Phase 2 studies had provided information on the effect of the drug on a surrogate outcome, no data were available on the primary endpoint of the key Phase 3 studies: moderate-to-severe asthma exacerbations, which are potentially life-threatening events with a significant burden on patients' lives<sup>19</sup>. Additionally, there was an important key secondary endpoint — forced expiratory volume in 1 second (FEV<sub>1</sub>), an endpoint commonly used in asthma trials —, for which Phase 2 data were available, but experts' judgements were sought on the effect of the different treatment duration and trial population in Phase 3. If the drug worked on one endpoint, it was considered to more likely work on the other endpoint. Thus, a joint distribution was required. Techniques to address both problems through expert elicitation are available within the SHELF framework.

In Section 3, we first give a brief overview of elicitation methods and of SHELF. Then we describe the extension method for eliciting judgements about Phase 3 outcomes by linking to Phase 2 results and the copula method for eliciting joint distributions. In Section 4, we return to our motivating example and describe how we used these techniques to estimate the PoS of that drug development program. We also compare the obtained expert judgements with the outcomes of the Phase 3 studies. We finish with some conclusions and recommendations in Section 5.

## 2 Motivating example

The Phase 3 program of fevipiprant, a prostaglandin D<sub>2</sub> receptor 2 antagonist for the treatment of asthma, was selected to pilot a new PoS framework that has since been introduced at Novartis<sup>20</sup>. At the time of the PoS assessment, fevipiprant had been studied in several Phase 2 randomised controlled trials and the Phase 3 clinical trials comparing two fevipiprant doses (150 or 450 mg once a day) with placebo were underway, with data collection almost complete. This

timing was one reason the program was selected as a pilot, because it ensured that the PoS assessment could not be influenced by the Phase 3 data, while at the same time minimising the time until the PoS assessment could be compared to the Phase 3 results. In reality, the assessment of the program and the decision to proceed with Phase 3 had already been taken at the end of Phase 2 based on more limited information.

One major challenge was that — unlike the Phase 2 trials — the key Phase 3 trials focused on more severe asthma patients with the sub-population with a blood eosinophil count  $\geq 250$  cells/ $\mu\text{l}$ . The primary null hypotheses for this sub-population were tested first in the trials’ testing procedures<sup>21</sup>. None of the Phase 2 trials evaluated the effect of fevipirant on moderate-to-severe asthma exacerbations. The annualised rate of such exacerbations was the primary endpoint of the two most important trials in the Phase 3 program<sup>21</sup>. Instead, a surrogate endpoint of reduction in sputum eosinophil counts had been measured in one of the Phase 2 trials<sup>22</sup>. FEV<sub>1</sub> was a key secondary endpoint in the Phase 3 program and has high regulatory acceptance as a measure of asthma control<sup>23</sup>. FEV<sub>1</sub> had been a primary or secondary endpoint of several of the Phase 2 studies including for dose ranging<sup>24</sup>, but these trials were of shorter duration and had a patient population with milder asthma than the Phase 3 trials.

As per the newly implemented PoS framework at Novartis, success was defined as regulatory approval with point estimates for key endpoints achieving or exceeding targets specified as part of a target product profile. It was assumed that regulatory approval would require statistical significance at the one-sided 0.025 significance level for at least one dose for both exacerbations and FEV<sub>1</sub> in both of the key Phase 3 trials. Thus, to calculate the PoS, we needed a joint prior distribution for the effects of fevipirant on exacerbations and FEV<sub>1</sub>.

Given the data that were available at the time of the PoS assessment, we decided to do this by eliciting the judgements of a group of experts. The question then was how best to structure the elicitation process: we wanted to explicitly leverage the Phase 2 data on the surrogate endpoint of reduction in sputum eosinophil counts, since this was arguably the most relevant evidence we had for informing beliefs about the effect of fevipirant on exacerbations.

We also expected experts to judge a larger effect of fevipirant on FEV<sub>1</sub> to be more likely the larger the effect of the drug on asthma exacerbations is. As a consequence, in order to fully characterise the joint distribution of these two treatment effects we would need to understand the size and direction of the dependence between these two quantities. In the next section, we describe the various approaches considered for the elicitation, before we return to the motivating example in Section 4 and describe how we practically applied these methods.

## 3 Overview of SHELF

### 3.1 Elicitation protocols

Elicitation can be done informally, but numerous pitfalls await the inexperienced practitioner, including well-established sources of bias in expert judgements<sup>25;9;26</sup>. Therefore, when the expert judgements are sufficiently important it is necessary to employ a formal procedure in the interests of quality and transparency. A small number of established elicitation *protocols* have been developed and refined by experienced practitioners. An overview is given by Dias et al.<sup>27</sup>.

### 3.2 The basic SHELF protocol and principles

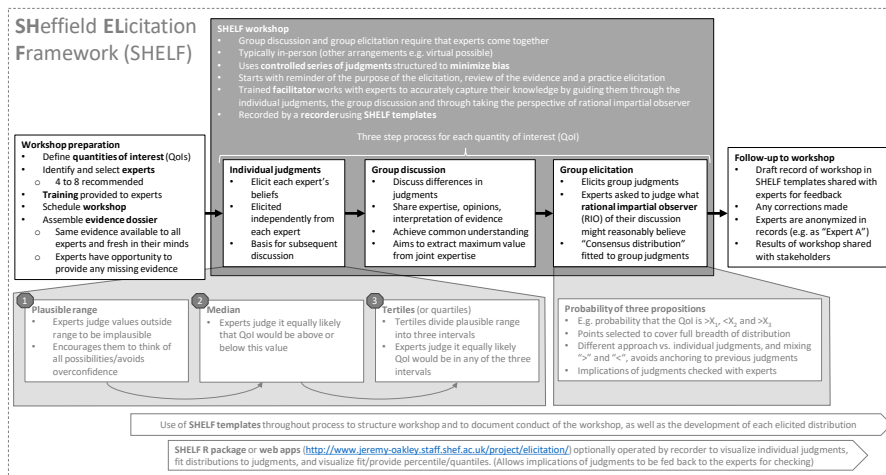
The SHELF protocol is characterised by carefully structured sequences of judgements designed to minimise biases and a unique way of eliciting a consensus probability distribution from a group of experts by combining individual and group elicitation<sup>28</sup>.

As shown in Figure 1, a workshop structured around the following three steps is central to the elicitation process:

1. Individual judgements: a plausible range of values, median and tertiles — or quartiles, but based on the authors’ experience tertiles are a reasonable default choice, because thinking about three instead of four equally likely intervals may be less challenging for experts — for the QoI are elicited from each expert independently.
2. Group discussion: The experts then share their beliefs and their rationale. This phase is where different interpretations and weighting of the available evidence are aired and debated.
3. Group elicitation: Judgements about the QoI (such as probabilities that the QoI is below one value, above another value or below a third value) are elicited, and a “consensus” distribution is fitted to these judgements.

Note that SHELF does not expect the experts to reach complete agreement such that they now have the same knowledge and beliefs about the QoI. Instead, they are asked to judge what a rational impartial observer, called RIO, might reasonably believe after seeing their individual judgements and listening to their discussion. By taking the perspective of RIO, experts can reach agreement on a distribution that represents a rational impartial view of their combined knowledge. The workshop facilitator has the important role to help the experts in accurately capturing their knowledge, facilitating the group discussion and leading them in applying the RIO perspective in order to avoid biases arising from the group interaction. The workshop recorder supports the facilitator through visualizing elicited judgments and taking minutes. At Novartis, an initial group of facilitators and recorders was selected and received a three-day training course on how to apply the SHELF protocol. Subsequently this group self-organised

Figure 1: Overview of the basic SHELF method



additional elicitation exercises and facilitators had the opportunity to observe a more experienced colleague before leading their first workshop.

Throughout the workshop the facilitator will prompt and challenge the experts to ensure that their statements genuinely represent what they believe or what RIO might believe. The facilitator will also work with the experts to identify a suitable distribution to represent their judgements about RIO's beliefs, which is not necessarily limited to the classes of distributions implemented in the SHELF R package<sup>29</sup>. The fitted distribution is the key outcome of the elicitation process.

The SHELF package freely available from the SHELF website<sup>30</sup> provides advice, templates and tools to support facilitators. It helps them to ask questions in such a way that biases are minimised and there is no need for the experts to have a thorough understanding of probability or statistical theory. Additionally, training in making these judgements is available for experts through an online self-paced course accessed from the SHELF website<sup>31</sup>. Since its inception in 2008, SHELF has been steadily expanded with new advice and methods. For example, the extension method described in Section 3.3 was added in version 4<sup>32</sup>.

### 3.3 The SHELF extension method

The package of SHELF materials<sup>32</sup> contains several techniques for eliciting a joint distribution for two or more uncertain quantities, including the extension and copula methods. The extension method is a generic technique that allows considerable flexibility for the form of the joint distribution. It is e.g. suitable for eliciting judgements about the treatment effect for a Phase 3 endpoint ( $X$ ) based on the Phase 2 results for a surrogate endpoint ( $Y$ ). The fact that Phase

3 follows Phase 2 chronologically makes it natural to express judgements about  $X$  conditional on  $Y$ .

For two QoIs,  $X$  and  $Y$ , the extension method consists of obtaining a marginal distribution for  $Y$  and a set of conditional distributions for  $X$  given  $Y = y$ . The elicitation of joint distributions requires the following steps, which are illustrated using the asthma example in Section 4.3.3.

1. A marginal distribution for  $Y$  is obtained. This distribution can be elicited as described in Section 3.2, but could also be the result of an analysis of available data. E.g. in the asthma case study introduced in Section 2 it is a meta-analytic predictive distribution<sup>33</sup> based on Phase 2 data.
2. A conditional distribution (as always, from the perspective of RIO) is elicited for  $X$  conditional on  $Y$  equalling the median of its elicited marginal distribution, also following the basic method of Section 3.2.
3. Several other quantiles of the elicited marginal distribution of  $Y$  are selected as conditioning points; typically these will be the quartiles, 5th and 95th percentiles. Median values are elicited for  $X$  conditional on  $Y$  equalling each of these conditioning points (first the 5th and 95th percentiles and then the quartiles). The basic SHELF approach of individual judgements – discussion – group judgements is used for each.
4. The final step is to ‘fit’ a set of conditional distributions to these judgements. First, a median function  $m(y)$  is fitted to the elicited conditional medians. This might for instance be a polynomial or a piecewise-linear fit (with extrapolation), and may be applied on a transformed scale. Second, a model is chosen to determine the conditional distributions based on the distribution at the  $Y$ -median elicited in Step 2. For instance, it may be decided that the  $Y$ -median distribution can be applied to all conditionals, simply shifted to follow the  $m(y)$  function. Alternatively, the variance may also be scaled depending on  $m(y)$ . These choices are available in the SHELF R software, but again other choices can be made. The facilitator will always work with the experts to identify a ‘fit’ that best represents their judgements.

The extension method is appropriate when the experts perceive a natural causal link from  $Y$  to  $X$ . Indeed, it is particularly useful when the objective is to elicit a distribution for  $X$  but the experts would find it easier to make judgements about  $X$  if they knew the value of  $Y$ . In this case, the marginal distribution of  $X$  is the main outcome of the elicitation process. Although it will not generally be feasible to derive that marginal distribution analytically from the elicited joint structure, a large Monte Carlo sample can be drawn by sampling values  $y_i$  from the marginal distribution of  $Y$  and then sampling  $x_i$  conditional on  $Y = y_i$ . The Monte Carlo samples  $\{x_i\}$  are then samples from the marginal distribution of  $X$  and, if needed, a distribution can be fitted to the samples.

### 3.4 The SHELF copula method

When there is no natural ordering of related QoIs based on time or causality, the extension method requires an arbitrary imposition of an ordering and the conditional judgements are more difficult for the experts. The SHELF copula method is appropriate for two or three QoIs and does not require the elicitation of conditional distributions. However, it does place some constraints on the joint distribution. The method has the following steps.

1. Marginal distributions are elicited for each QoI individually, using the basic method of Section 3.2.
2. For each pair of QoIs, a single judgement concerning their degree of correlation is made. This judgement is called the concordance probability, and is the probability that both QoIs lie on the same side of their respective elicited medians<sup>34</sup>.
3. A Gaussian copula joint distribution<sup>35</sup> is then fitted to these marginal distributions and concordance probabilities. The facilitator shows the experts suitable displays or summaries of the joint distribution to verify that it is a reasonable representation of their beliefs.

With just two QoIs, the copula method is simple to apply. The Gaussian copula imposes a restriction on the joint distribution but in practice it will usually be an adequate fit to the experts' judgements. The SHELF R package makes it straightforward to carry out these steps, to draw samples from the joint distribution using the `copulaSample` function and to explore the various choices to be made. Full technical details and advice are available in the package of SHELF materials<sup>32</sup> for those users who may wish to implement the method in some other software or to obtain a deeper understanding of it.

In principle, the copula method is applicable for larger numbers of QoIs, but it is difficult to use for more than three for the following reasons. With three QoIs, three concordance probabilities need to be elicited. Under the Gaussian copula assumption, each concordance probability  $p_c$  can be transformed to a correlation coefficient  $\rho = \sin(\pi(p_c - 0.5))$ <sup>34</sup> and the resulting correlation matrix must be positive definite. It is quite possible for the experts' elicited concordance probabilities to fail to produce a valid correlation matrix, and they must then revisit their judgements with the aid of the facilitator to achieve an adequate fit. With more than three QoIs, the number of concordance probabilities that needs to be elicited rapidly increases, as does the likelihood of the elicited values not corresponding to a valid correlation matrix.

The SHELF copula method is a natural choice to construct a joint distribution for the effects of a drug on two Phase 3 endpoints, such as a primary and secondary clinical outcome.

The interested reader will find full technical details, as well as much practical advice, on these and other elicitation techniques in the package of SHELF materials<sup>32</sup>.



## 4 Asthma case study

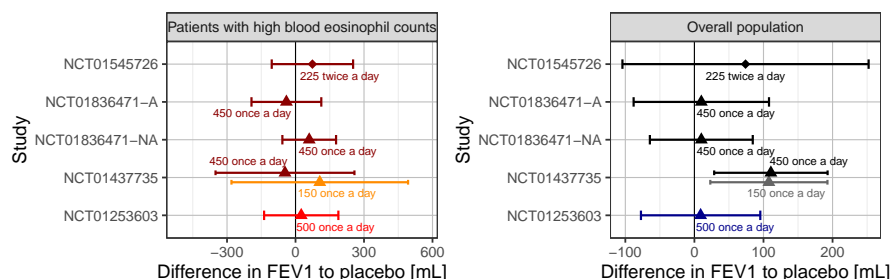
In this section, we provide an in-depth description of the expert elicitation and PoS calculation for the example introduced in Section 2. We decided to structure the elicitation process into three parts. First, we followed the SHELF extension method by using Phase 2 data to establish a marginal distribution for the effect of fevipirant on sputum eosinophil counts and then elicited from a group of experts a set of conditional judgements on the effect on exacerbations in the Phase 3 population given different values for the effects on this surrogate endpoint. Secondly, we elicited the experts' beliefs on the effect of fevipirant on FEV<sub>1</sub> in the Phase 3 population. Finally, we used the SHELF copula method to elicit the dependence between drug effects on exacerbations and FEV<sub>1</sub>.

### 4.1 Available evidence

Fevipirant was studied in four Phase 2 randomised controlled trials in asthma and the results of these studies for the FEV<sub>1</sub> endpoint are summarised in Figure 2.


1. A Proof of Concept trial (ClinicalTrials.gov identifier NCT01253603) with a 4 week treatment duration in patients on reliever therapy did not show an effect of fevipirant on the primary endpoint of FEV<sub>1</sub> in the overall trial population, but more favourable results were seen for a subgroup of more severe patients<sup>36</sup>.
2. A dose finding trial (NCT01437735) with a 12 week treatment duration<sup>24</sup> was the basis of the selection of one of the Phase 3 doses.
3. A 12-week trial looked at potential differences in effects in patients with atopic and non-atopic asthma (NCT01836471).

Figure 2: Observed differences in forced expiratory volume in 1 second (FEV<sub>1</sub>) to placebo with 95% confidence intervals for fevipirant in Phase 2 studies in the subgroup with a blood eosinophil count of  $\geq 250$  cells/ $\mu$ L and in the overall trial populations (A: atopic patients, NA: non-atopic patients)



- Finally, there was a trial (NCT01545726) that showed a reduction of sputum eosinophil counts after 12 weeks of treatment with fevipirant compared with a placebo group<sup>22</sup>. The ratio of a 3.5-fold (95% CI 1.7 to 7.0;  $p=0.0014$ ) lower ratio of geometric means in sputum eosinophil counts from baseline to the end of treatment compared with placebo in this study was a key rationale for initiating Phase 3 trials investigating an effect on asthma exacerbations<sup>37</sup>. Figure 3 shows these trial results, as well as a predictive distribution for the true value of this ratio in a new study given the results of this study. The prior used for the predictive distribution was based on an industry benchmark as described in Section 4.4<sup>20</sup>.

Figure 3: Point estimate and 95% confidence interval, and predictive distribution with median and 95% prediction interval for the mean in a new study

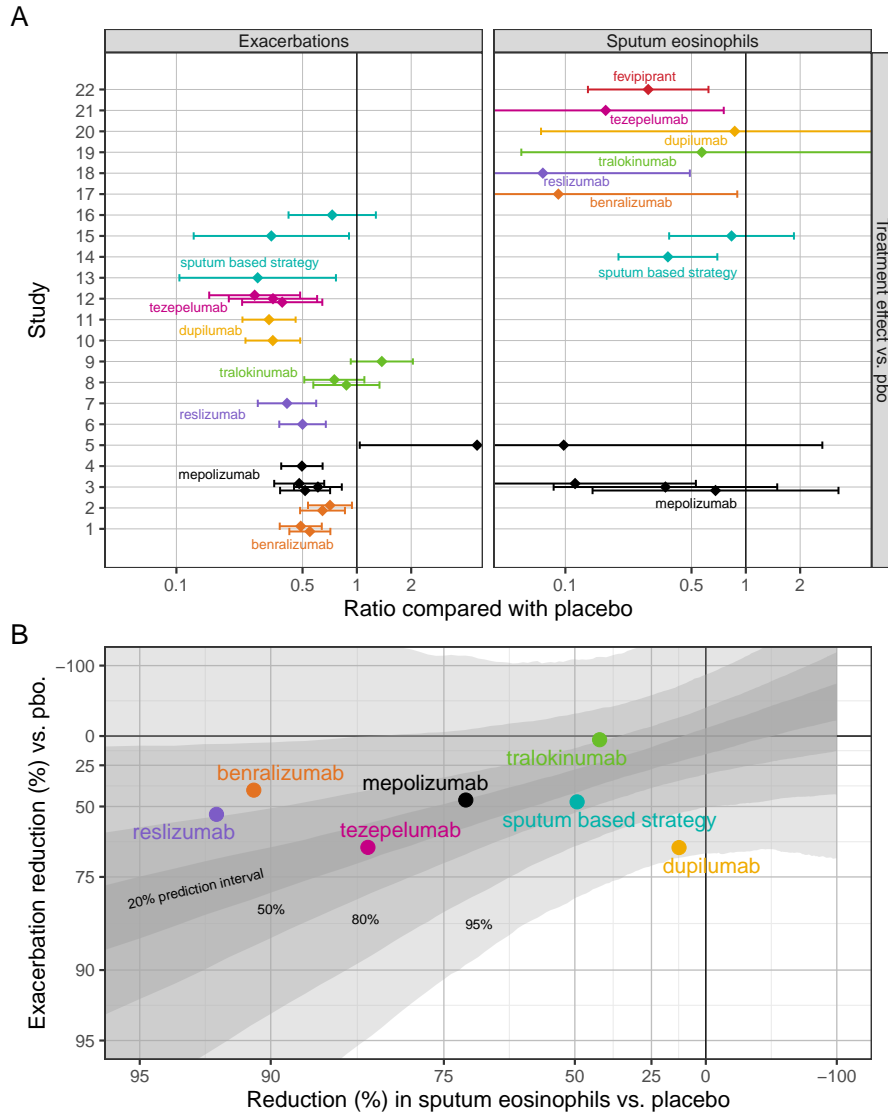


A number of anti-inflammatory treatments that lower sputum eosinophil counts have been shown to reduce exacerbation rates in asthma patients with elevated sputum eosinophil counts<sup>38</sup>. This evidence was mostly generated with corticosteroids, but suggests that sputum eosinophil counts may be a surrogate for a reduction in exacerbations. As part of the evidence dossier for this expert elicitation, we assembled more recent evidence from 22 trials of other drug classes<sup>39;40;41;42;43;44;45;46;47;48;49;50;51;52;53;54;55;56</sup>. The data are shown in Panel A of Figure 4 and the results from a Bayesian meta-regression model are shown in Panel B of the figure. Without data from a variety of different drugs, this meta-regression would be highly questionable, because then its findings might only apply to a specific mode of action, i.e. the specific way a drug produces an effect in the body. Note that some of these data were not available at the time the Phase 3 program for fevipirant was started.

For the question of the likely effect of fevipirant on  $FEV_1$  in asthma patients with blood eosinophil counts  $\geq 250$  cells/ $\mu$ L, the evidence dossier presented the Phase 2 results for the overall population, as well as for subgroups defined by blood eosinophil counts (see Figure 2). In the overall population, only study NCT01437735 achieved a statistically significant superiority compared with placebo, while for the exploratory analyses by blood eosinophil count, all confidence intervals included no effect of fevipirant over placebo.

In addition, the evidence dossier gave details of the fevipirant Phase 3 program, and discussed the strengths and limitations of the available evidence

Figure 4: Effects of anti-inflammatory asthma therapies on sputum eosinophil counts and exacerbation rates compared with placebo: Estimates with 95% confidence intervals for exacerbation rate ratios and ratio of geometric mean (vs. placebo) ratios of sputum eosinophil levels at the end of the study compared with baseline (Panel A), and meta-regression using random drug effects on intercept and slope of relationship, as well as random study effects (Panel B); Studies 10 and 11 are the two parts of study NCT02414854 that were not blinded against each other.



that the experts needed to bear in mind.

## 4.2 Choice of quantities of interest for elicitation

The QoI to be elicited were chosen based on their importance for meeting the success definition of the PoS framework and lack of evidence to directly inform a predictive distribution. The global project team considered the results in the two exacerbation trials (NCT02555683 and NCT02563067) in the pre-specified subgroup of patients with high eosinophil counts to be the most important to fulfil the target product profile. These 1-year exacerbation trials compared two doses of fevipiprant with a placebo on top of continued standard of care therapy in severe asthma patients. The rate of asthma exacerbations (target product profile criterion:  $\geq 40\%$  relative rate reduction compared with placebo) was the primary endpoint of these studies, while the key secondary endpoint of FEV<sub>1</sub> (target product profile criterion:  $\geq 120$  mL improvement in FEV<sub>1</sub> compared with placebo) was considered to be especially important for regulatory approval. There was considerable historical data on the placebo exacerbation rate, the between patient heterogeneity in the exacerbation rate<sup>57</sup> and the variability in FEV<sub>1</sub> so that these quantities did not require elicitation.

The biggest source of uncertainty regarding the PoS was about the effects of fevipiprant on asthma exacerbations and FEV<sub>1</sub>, as well as about their correlation. For this reason, these were identified as the QoIs for the expert elicitation. We carefully chose the phrasing of the questions about the QoIs to make it easy for the experts to think about them and express their judgements.

We decided to use the extension method to elicit judgements about the relative rate reduction in exacerbations conditional on a specified reduction in sputum eosinophils, and to use the copula method to elicit the association between the two QoIs. On that basis, we formally defined the following three QoIs:

- $X$  is the average reduction in moderate to severe asthma exacerbations achieved by fevipiprant compared to placebo over the population of eligible patients,
- $Y$  is the average reduction in sputum eosinophil counts achieved by fevipiprant compared to placebo over the population of eligible patients,
- $Z$  is the average increase in FEV<sub>1</sub> achieved by fevipiprant compared to placebo over the population of eligible patients.

Eligible patients are defined as matching the inclusion criteria for the NCT02555683 and NCT02563067 Phase 3 trials and having blood eosinophil counts of at least 250 cells/ $\mu$ L. Note that because we had already derived the marginal predictive distribution in Figure 3 for the reduction  $Y$  in sputum eosinophil counts from Phase 2 data, the extension method for the QoI  $X$  required only conditional distributions to be elicited.

The choice and phrasing of the QoIs in elicitation is an important early task. Quantities must be clearly and unambiguously defined, in terms that are

familiar to the experts. It must be clear that each quantity has a unique, well-defined (but unknown) value. We chose to elicit treatment effects compared with placebo as percentage reductions in exacerbations and improvements in FEV<sub>1</sub>, because these are widely used effect measures in asthma trials commonly expressed in these terms that were familiar to the experts. The effects are defined as averages over all potential patients so that they have well-defined and unique values. The experts would be asked for their judgements on questions such as:

1. Given that an anti-inflammatory drug reduces sputum eosinophil counts by  $Y$ , what do you judge to be the likely values for the relative exacerbation rate reduction  $X$  in eligible patients?
2. What do you judge to be the likely values for the difference  $Z$  between fevipirant and placebo in FEV<sub>1</sub> in millilitres (mL) in eligible patients?
3. Given the judgements about the reduction in exacerbations and the change in FEV<sub>1</sub> caused by fevipirant, how likely do you judge it to be that both  $Y$  and  $Z$  will be on the same side of your median values?

### **4.3 The elicitation workshop**

#### **4.3.1 Experts for the elicitation workshop**

In order to capture the full range of opinions and differing past experiences amongst experts, a group of Novartis internal experts was convened. The 5 selected experts all had extensive experience in drug development in the respiratory area. Two were part of the fevipirant team (a clinician and a statistician), while 3 were not members of the fevipirant team (a clinician, a translational medicine expert and a regulatory affairs expert). These experts were selected, because the QoIs appeared to be related to clinical trials and understanding mechanistic considerations around the drug efficacy. We wanted at least some of this key expertise to be from outside of the fevipirant project team to ensure an outside opinion would be heard. A statistician was considered important to provide a perspective on the available evidence and the expert in regulatory affairs was selected due to a broad experience with multiple previous programs.

Prior to the elicitation workshop, all experts were encouraged to work through an online course on expert elicitation<sup>31</sup> and they were guided through a practice exercise by the facilitator at the start of the workshop.

#### **4.3.2 Conduct of the elicitation workshop**

The elicitation workshop was an in-person 4-hour meeting with one facilitator, one recorder and five experts. While the facilitator guided the meeting and asked the experts questions, the role of the recorder was to operate the SHELF software, project relevant visualisations for the experts and to take minutes of the meeting.

### 4.3.3 Elicitation of first quantity of interest

The median of the marginal distribution of  $Y$  shown in Figure 3 — based on a Bayesian analysis of Phase 2 sputum eosinophil data — was a 66% reduction (80% interval from 52 to 76%). Round numbers are easier for experts to condition on, and so, for the first QoI, the median of 66% was rounded to 65%. Thus, the experts were first asked for their judgement on  $Y$  conditional on  $X$  being a 65% reduction in sputum eosinophil counts.

For the individual judgements about this QoI, the tertile method was used. Each expert first independently wrote down their plausible range for the QoI, followed by their median and the points that divide the plausible range into equally probable thirds. At each step the experts were asked to challenge their own judgements. For instance, after specifying their plausible range, experts were asked to consider their reaction if a large study estimated  $X$  to be outside that range; would they acknowledge that their range was too narrow, or would they be suspicious of the reported estimate? If their reaction would be the former one, then they should widen their plausible range.

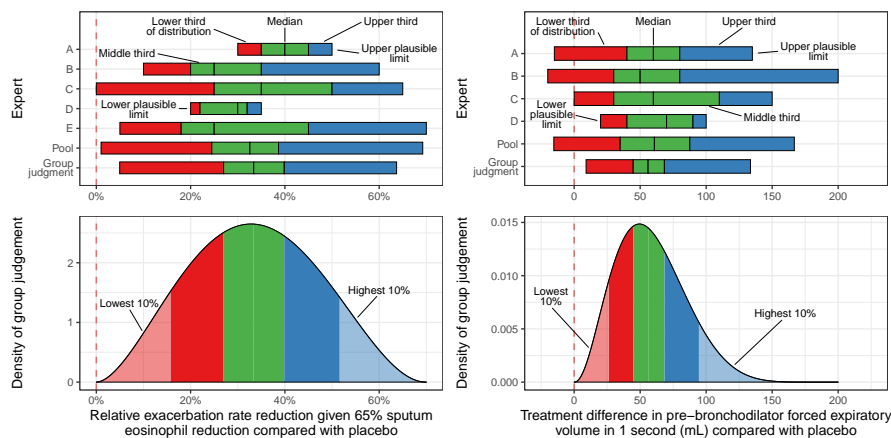
Then the individual judgements were revealed to the group and the experts were asked to explain their judgements. In this wide-ranging discussion, a number of points were raised and the main arguments were recorded using the SHELF templates. Afterwards, consensus judgements were obtained using the probability method: experts were asked what probability RIO (the Rational Impartial Observer) would assign to the relative exacerbation rate reduction being less than 25%, greater than 40% and less than 35%. After significant discussion, the group agreed that RIO would assign probabilities of 30%, 30% and 50%, respectively. The values of 25%, 40% and 35% were chosen to characterise the lower, upper and middle part of the consensus distribution. Note that these values lie near the lower and upper tertile, and near the median of a linear pool of the opinions of the individual experts. Switching between "less than" and "greater than" is intended to reduce anchoring of each judgement to the previous judgement. The order of the questions aims to make the experts first think of the full range of plausible values, before eliciting the location of the center of the distribution.

A Beta(2.81, 3.05) distribution scaled to a plausible range of 0 to 70% was fitted to these judgements and shown to the experts. The experts felt that this distribution, with a median at 33.4% (90% credible interval 11.9 to 55.8%), adequately represented their knowledge.

The result of this elicitation was a distribution for  $X$  (exacerbation reduction), given that  $Y$  (sputum eosinophil reduction) is 65%. The results of the individual judgements and the group judgement are shown on the left-hand side of Figure 5.

Then, the experts were asked for their conditional judgement about the median percentage reduction in exacerbations given an effect on sputum eosinophil of 50%, then for 75%, 60% and 70%. These numbers correspond approximately to 10%, 90%, 25% and 75% points of the marginal predictive distribution for effects of fevipiprant on sputum eosinophil counts, respectively. Thus, they

Figure 5: Distributions elicited from individual experts, linear pool of these distributions and group judgements



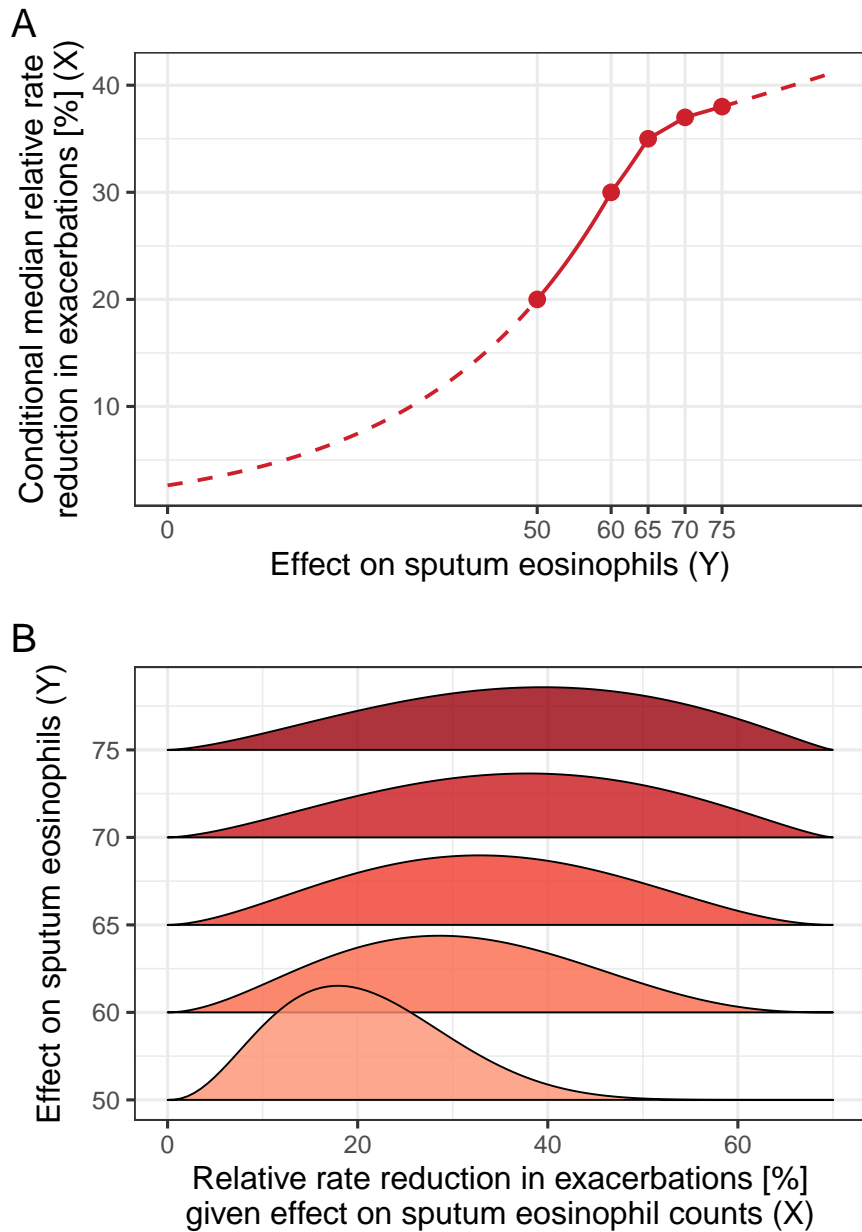
characterise conditional judgements across the bulk of this distribution. Their order was chosen in order to minimise known sources of cognitive bias and to ensure that experts needed to think carefully about each judgement. The elicited medians are shown in Panel A of Figure 6.

It was agreed that over the plausible range of effects on sputum eosinophil counts, there was no probability that the drug could increase the number of exacerbations, because the assumption that fevipiprant reduced sputum eosinophils indicated at least some positive benefit. It was therefore appropriate to model the distributions of exacerbation reductions at intermediate sputum eosinophil effects through a log transformation — i.e. to assume that  $\text{median}(\log(X|Y))$  is a piecewise linear function of  $Y$ . The experts were shown the resulting median relationship shown in Panel A of Figure 6 and agreed that it represented a reasonable RIO opinion.

Using the log transformation, the conditional distribution given  $Y = 65\%$  was assumed for  $X$  conditional on other values of  $Y$ , but scaled to follow the elicited median model — i.e. we shifted the median of each Beta-distribution according to Panel A of the figure and kept the variance on the log-scale constant. The recorder showed the experts the resulting conditional distribution plot in Panel B of Figure 6. The facilitator pointed out how the scaling had resulted in less uncertainty conditional on  $Y = 50\%$  but more conditional on  $Y = 75\%$ . The experts confirmed that this was a reasonable representation of their beliefs.

The elicitation of the first QoI was now complete and the required (marginal) distribution for  $X$  was computed by Monte Carlo simulation by combining the elicited conditional relationship with the predictive distribution for  $Y$  from Figure 3. It is shown in the top-most panel of Figure 8.

Figure 6: Piecewise-linear median model for the elicited medians (Panel A) and conditional distributions for the relative exacerbation rate reduction across the range of plausible effects on sputum eosinophil counts (Panel B)





#### 4.3.4 Elicitation of further quantities of interest

The elicitation for the second QoI then proceeded using the tertile method for individual judgements, followed by a discussion and, again, using the probability method for the consensus judgement. The resulting judgements are shown on the right-hand side of Figure 5.

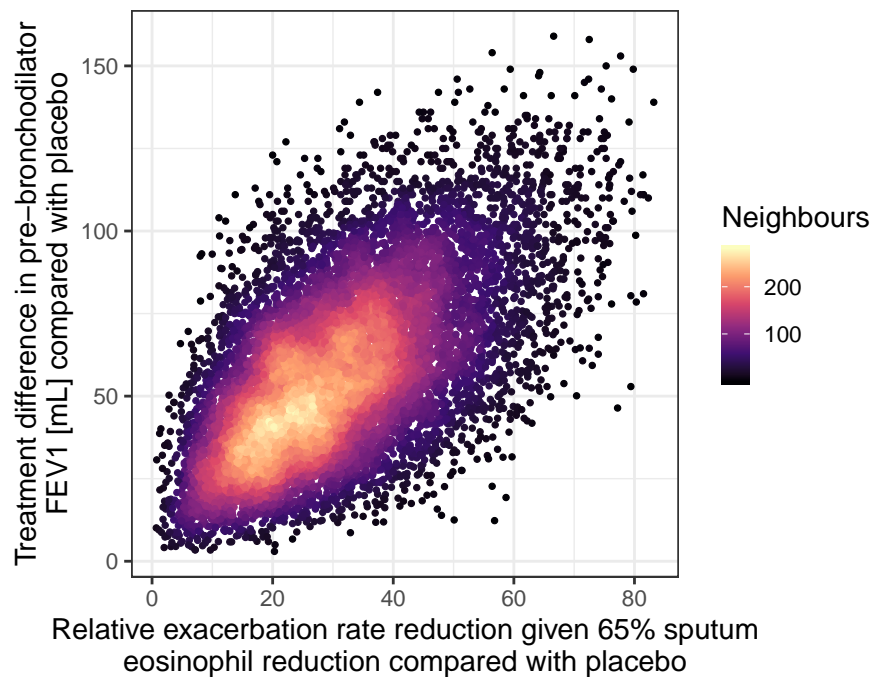
The joint distribution of the treatment effects on exacerbations and FEV<sub>1</sub>,  $X$  and  $Y$ , was then elicited using the copula method. The correlation was elicited through the concordance probability, i.e. RIO's judgement of the probability that the true values of  $X$  and  $Y$  would both be on the same side of their elicited medians. The experts were shown a figure with four quadrants (above and below the medians for exacerbations and FEV<sub>1</sub>) clearly marked, which was used as a visual aid. The experts found the concordance probability difficult to judge. After the facilitator gave an alternative explanation in terms of the conditional probability that one variable was above its median given that the other was above its median, a concordance probability of 0.7 was tentatively agreed by the experts. The experts were shown a graphic similar to Figure 7 for the case of a concordance probability of 0.7 and found it very helpful and in accord with their expectations. Alternative concordance probabilities were explored using the same graphical display. The correlation was too tight with 0.8 concordance and the experts felt that there was appreciable positive correlation so 0.5 concordance was not considered appropriate. It was important to elicit the correlation between the treatment effects on exacerbations and FEV<sub>1</sub>, because a higher concordance probability increased the PoS and a sensitivity analysis with concordance probabilities between 0.5 (correlation 0) to 0.856 (correlation 0.9) changed the odds of success by a factor of up to 2.

The resulting joint distribution is shown in Figure 7 and constitutes the joint predictive prior distribution for the true treatment effect on the two outcomes in the Phase 3 trials implied by the judgements elicited from the experts.

#### 4.4 Probability of success calculation

We already described the basic aims of the newly introduced PoS framework at Novartis at a high level in Section 2. Its practical application involves the following four steps<sup>20</sup>. First, a benchmark probability of approval for a project at the start of Phase 2 is estimated based on a small number of program characteristics by a logistic regression model trained on a database of drug development projects. Second, a Bayesian analysis is conducted, in which the prior for the efficacy effects is set based on the benchmark probability of efficacy success in both Phase 2 and 3. This prior is then used in combination with Phase 2 data to obtain a posterior distribution for drug efficacy. Phase 3 studies are then simulated using samples from the posterior in order to estimate the probability of the key efficacy endpoints meeting target product profile criteria in the Phase 3 program. Benchmark information is also used to account for the risk of program failure due to an unexpected safety issue and of not obtaining regulatory approval despite a successful Phase 3 program. Third, a program risk assessment

Figure 7: Point density plot of elicited joint distribution for treatment effect on exacerbations and forced expiratory volume in 1 second ( $FEV_1$ ) based on 10,000 Monte Carlo samples



is done to capture other risks not already covered by the previous calculations. This assessment is then used to adjust the probability of a registration with a label meeting target product profile criteria to obtain the PoS. The adjustment in this step was also determined using elicitation process. Finally, in exceptional circumstances a fourth step allows for an adjustment for factors not captured by the preceding three steps.

In this case study, the Bayesian analysis in the second step of the PoS approach could not directly inform the PoS of the Phase 3 program due to the differences in endpoints and population between Phase 2 and 3. Thus, the results of the Bayesian analysis for sputum eosinophil counts in Figure 3 were linked to the efficacy on asthma exacerbations in Phase 3 using an expert elicitation in the manner described in Section 3.3. In contrast, the effect of fevipiprant on FEV<sub>1</sub> was elicited directly from the experts and the joint distribution of the efficacy of fevipiprant for both endpoints was then obtained as described in Section 3.4.

For pragmatic reasons the Novartis PoS approach foresees that only one or two key endpoints should be considered in the definition of success. For this reason, it was decided to ignore the other two key secondary endpoints (asthma control questionnaire and asthma related quality of life questionnaire) of these Phase 3 trials for the purposes of the PoS calculation.

#### 4.4.1 Calculation of PoS estimates

The estimated benchmarks for the first indication of a respiratory orally administered small molecule targeting a receptor:

- a Phase 2b efficacy success probability of 54% conditional on starting Phase 2b,
- a Phase 3 efficacy success probability of 68% conditional on Phase 2b success, and
- an approval probability of 94% conditional on Phase 2b and 3 success.

The program risk assessment<sup>20</sup> considered the majority of categories to fall into the lowest risk category with one question falling into the intermediate risk category.

Phase 2b and 3 benchmark efficacy probabilities were used to define the prior distribution used to obtain the predictive distribution for the drug effect on sputum eosinophils shown in Figure 3 as described by Hampson et al.<sup>20</sup>. This served as an input to the expert elicitation as described in Phase 3 outcomes were simulated assuming the actual trial design and sample size, the protocol assumptions for nuisance parameters, and the prior predictive distribution for the treatment effect shown in Figure 7. 10% of the simulated trials achieved statistical significance for both exacerbations and FEV<sub>1</sub> for one dose in both trials, and in 4% of simulations the point estimates achieved the target product profile criteria. This probability is multiplied with the probability of no safety

showstopper in Phase 3 of 94%, and the probability of approval and meeting additional market access hurdles conditional on Phase 3 success of 93%. This resulted in a final PoS of 4% after rounding.

The main hurdle was identified to be the FEV<sub>1</sub> endpoint and the ambitious target product profile for exacerbations. If one only considered a target product profile requiring a relative exacerbation reduction of 30% with no requirements for FEV<sub>1</sub>, the PoS became 46%.

## 4.5 Timelines

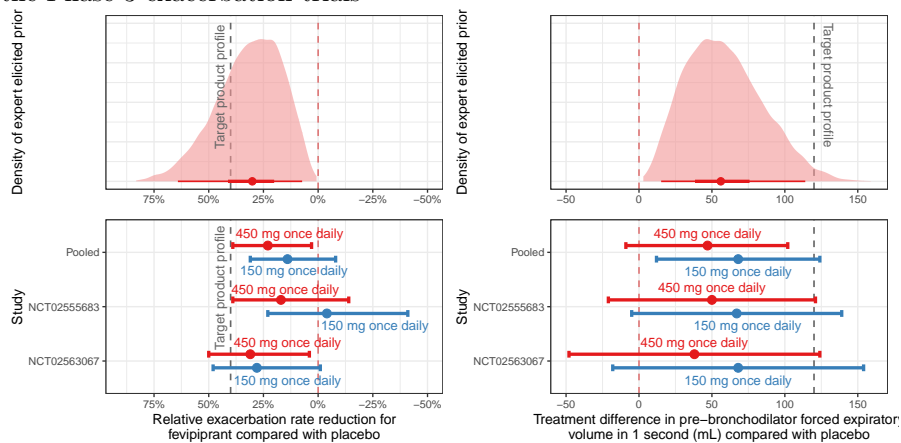
The whole PoS process required approximately 2 months. After an initial review, we identified that an expert elicitation workshop would be needed. On 28 May 2019, we identified the facilitator for the workshop and compiled a list of candidate dates. In the meantime, the team worked to assemble an evidence dossier. By 12 June, we had arranged an elicitation workshop on 12 July after confirming the availability of five experts. By 1 July, the evidence dossier had been drafted by the biostatistics team, was shared with the facilitator and recorder, and was finalised on 8 July after a review by internal experts, four days before the workshop. One learning was that we should have shared the dossier with the experts earlier in order to allow them to provide feedback on its contents so that additional evidence could have been introduced up-front. On 12 July the workshop took place using version 4 of the SHELF methodology and on 20 July 2019 the final report of the elicitation meeting was issued. All recordings from the meeting were made using the templates provided as part of the SHELF documents package and participants were kept anonymous in these minutes by using the letters A to E for the experts, as well as Z for the facilitator.

## 4.6 Phase 3 results

The results of the Phase 3 trials, for which we conducted the expert elicitation, are shown in Figure 8. As can be seen only one comparison within one of the two trials was associated with a confidence interval that excluded no effect, but this result was not considered statistically significant after an adjustment for multiplicity<sup>21</sup>. The results of the Phase 3 trials are very informative in the sense that the 95% confidence intervals essentially exclude the target effect size of the target product profile .

These results are consistent with the elicited prior information from the experts: the experts essentially excluded the possibility that the true effect of the studied fevipirant doses on FEV<sub>1</sub> would meet the target product profile criteria, while for the primary exacerbation endpoint, the experts judged that there was a reasonable possibility that the true effect was at or above the target product profile requirement. On the basis of these Phase 3 results Novartis did not pursue a filing for an indication in asthma.

Figure 8: Implied distribution for true effect of fevipiprant 450 mg once daily on exacerbations and forced expiratory volume in 1 second ( $FEV_1$ ) based on elicited expert judgements, and study results in the high blood eosinophil subgroup of the Phase 3 exacerbation trials



## 5 Discussion

The quality of decisions in the presence of uncertainty can be improved by taking the judgements of experts based on the available evidence into account. When stakes are high, as with major investment decisions by a pharmaceutical company, the necessary effort and cost of obtaining experts' judgements is negligible compared to the cost of a wrong decision. This is one of the reasons why the new Novartis PoS framework, which is applied for the decision to initiate pivotal trials for a project, recommends expert elicitation when substantial direct evidence about QoIs is not available. The SHELF extension method and the SHELF copula method address two common scenarios in this setting: when we extrapolate the evidence from surrogate endpoints to Phase 3 endpoints, and when how much a drug affects one endpoint changes how much we judge it to affect other endpoints.

There are currently no published examples of how to apply these methods as part of the SHELF protocol in the pharmaceutical industry. Therefore, we felt it would be helpful to share an example illustrating the full extent of real-world complexities and the relevant practical considerations. This will hopefully help others that wish to use expert elicitation to inform clinical drug development or other types of high stakes decisions.

We do not wish to overemphasise the outcomes from a single example. Nevertheless, the close alignment between the experts' group judgements with the trial outcomes, which were not known to the experts at the time of the elicitation workshop, supports the validity of expert elicitation in drug development. If a similar elicitation outcome had been available at the time of the decision to

start the Phase 3 program for fevipirant, it would have suggested a lower PoS than assigned at the time and may have led to re-evaluation of the assumptions regarding the secondary FEV<sub>1</sub> endpoint. However, this proof of concept for elicitation as part of a new PoS framework was performed 4 years after this decision that was known to our experts and used information that only became available subsequently.

The project team noted that the evidence dossier and the discussions in the elicitation workshop were extremely helpful for assembling and understanding the existing evidence on the efficacy of the drug. It may sometimes be the case that teams are very well aware of the clinical trials conducted for their product, but have not systematically reviewed the indirect evidence that is available from other sources. After the elicitation workshop the experts expressed that they appreciated the structured and scientific process, that they found the methodology intuitive, and that they were positively surprised how fully non-statisticians could participate in the workshop.

While we describe a particular example of an elicitation workshop, we have now run several similar workshops at Novartis and some of the authors of this paper have several years of experience of doing so with other clients. On this basis, we offer a number of practical recommendations. It is important to start preparing the evidence dossier as early as possible so that experts and other stakeholders can give feedback prior to a workshop. This is also an opportunity to let senior leaders with strong positive opinions on projects provide the evidence they wish to be considered. Additionally, it can be difficult for experts to free their agenda for long workshops and we have found that people find it hard to concentrate in virtual meetings for as long as in in-person workshops. This has led us to investigate options for eliciting individual judgements prior to the main workshop, but this may reduce the quality of judgements obtained<sup>58</sup>. Thus, it is still important to conduct the group discussion and group elicitation part of the SHELF protocol with all experts in the same (physical or virtual) meeting, which provides ample opportunity to correct any ill-considered judgements.

It is also important to clearly communicate how elicitation results will be used. In the context of the PoS of drug development programs, this meant making it clear that the resulting probability is not the sole determinant of funding for a project. We now routinely remind teams that investment decisions will also be based on other factors such as the costs of development, market opportunity and unmet medical need.

## Acknowledgments

We thank Ana-Maria Tanase, Christian Hasenfratz and Hanns-Christian Tillmann for being experts for the asthma case study, as well as Kelvin Stott, Giovanni Della Cioppa and Karine Baudou for their support of the pilot phase of the Novartis PoS initiative.

## **Data Availability Statement**

The data that support the case study are all graphically displayed as part of the main article.

## References

1. Garthwaite Paul H, O'Hagan Anthony. Quantifying expert opinion in the UK water industry: an experimental study *Journal of the Royal Statistical Society: Series D (The Statistician)*. 2000;49:455–477.
2. Gosling John Paul, Hart Andy, Mouat David C, Sabirovic Mirzet, Scanlan Simon, Simmons Alick. Quantifying experts' uncertainty about the future cost of exotic diseases *Risk Analysis: An International Journal*. 2012;32:881–893.
3. Usher Will, Strachan Neil. An expert elicitation of climate, energy and economic uncertainties *Energy policy*. 2013;61:811–821.
4. Bamber Jonathan L, Oppenheimer Michael, Kopp Robert E, Aspinall Willy P, Cooke Roger M. Ice sheet contributions to future sea-level rise from structured expert judgment *Proceedings of the National Academy of Sciences*. 2019;116:11195–11200.
5. Kinnersley Nelson, Day Simon. Structured approach to the elicitation of expert beliefs for a Bayesian-designed clinical trial: a case study *Pharmaceutical statistics*. 2013;12:104–113.
6. Dallow Nigel, Best Nicky, Montague Timothy H. Better decision making in drug development through adoption of formal prior elicitation *Pharmaceutical Statistics*. 2018;17:301–316.
7. Ren S, Oakley JE, Stevens JW. Evidence Synthesis For Health Technology Assessment With Limited Studies *Value in Health*. 2017;20:A770.
8. Soares Marta O, Bojke Laura. Expert elicitation to inform health technology assessment in *Elicitation:479–494*Springer 2018.
9. European Food Safety Authority . Guidance on expert knowledge elicitation in food and feed safety risk assessment *EFSA Journal*. 2014;12:3734.
10. Daneshkhah Alireza, Oakley JE. Eliciting multivariate probability distributions *Rethinking risk measurement and reporting*. 2010;1:23.
11. Werner Christoph, Hanea Anca M, Morales-Nápoles Oswaldo. Eliciting multivariate uncertainty from experts: Considerations and approaches along the expert judgement process in *Elicitation:171–210*Springer 2018.
12. Elfadaly Fadlalla G, Garthwaite Paul H. Eliciting Dirichlet and Connor–Mosimann prior distributions for multinomial models *Test*. 2013;22:628–646.
13. Zapata-Vázquez Rita Esther, O'Hagan Anthony, Soares Bastos Leonardo. Eliciting expert judgements about a set of proportions *Journal of Applied Statistics*. 2014;41:1919–1933.



14. Baey Charlotte, Sahlin Ullrika, Clough Yann, Smith Henrik G. A model to account for data dependency when estimating floral cover in different land use types over a season *Environmental and ecological statistics*. 2017;24:505–527.
15. Norrington Lisa, Quigley John, Russell Ashley, Meer Robert. Modelling the reliability of search and rescue operations with Bayesian Belief Networks *Reliability Engineering & System Safety*. 2008;93:940–949.
16. Sigurdsson JH, Walls LA, Quigley JL. Bayesian belief nets for managing expert judgement and modelling reliability *Quality and Reliability Engineering International*. 2001;17:181–190.
17. Bedford Tim, Cooke Roger M. Probability density decomposition for conditionally dependent random variables modeled by vines *Annals of Mathematics and Artificial intelligence*. 2001;32:245–268.
18. Truong Phuong N, Heuvelink Gerard BM, Gosling John Paul. Web-based tool for expert elicitation of the variogram *Computers & geosciences*. 2013;51:390–399.
19. Global Initiative for Asthma . Global Strategy for Asthma Management and Prevention 2020. Available at <https://ginasthma.org/>.
20. Hampson Lisa V., Bornkamp Björn, Holzhauser Björn, et al. Improving the assessment of the probability of success in late stage drug development *Pharmaceutical Statistics*. 2022;21:439-459.
21. Brightling Christopher E, Gaga Mina, Inoue Hiromasa, et al. Effectiveness of fevipiprant in reducing exacerbations in patients with severe asthma (LUSTER-1 and LUSTER-2): two phase 3 randomised controlled trials *The Lancet Respiratory Medicine*. 2020.
22. Gonem Sherif, Berair Rachid, Singapuri Amisha, et al. Fevipiprant, a prostaglandin D<sub>2</sub> receptor 2 antagonist, in patients with persistent eosinophilic asthma: a single-centre, randomised, double-blind, parallel-group, placebo-controlled trial *The Lancet Respiratory Medicine*. 2016;4:699–707.
23. Committee for Medicinal Products for Human Use . Guideline on the clinical investigation of medicinal products for the treatment of asthma 2015. CHMP/EWP/2922/01 Rev.1.
24. Bateman Eric D., Guerrerros Alfredo G., Brockhaus Florian, et al. Fevipiprant, an oral prostaglandin DP<sub>2</sub>receptor (CRTh<sub>2</sub>) antagonist, in allergic asthma uncontrolled on low-dose inhaled corticosteroids *European Respiratory Journal*. 2017;50:1700670.

25. O'Hagan Anthony, Buck Caitlin E, Daneshkhah Alireza, et al. *Uncertain judgements: eliciting experts' probabilities*. Chichester: John Wiley & Sons 2006.
26. O'Hagan Anthony. Expert knowledge elicitation: subjective but scientific *The American Statistician*. 2019;73:69–81.
27. Dias Luis C, Morton Alec, Quigley John. Elicitation *Springer International Publishing*. MR3700912. doi: <https://doi.org/10.1007/978-3-319-65052-4>. 2018;1:3.
28. Gosling John Paul. SHELF: the Sheffield elicitation framework in *Elicitation*:61–93Springer 2018.
29. Oakley Jeremy. *SHELF - Tools to Support the Sheffield Elicitation Framework* 2020. R package version 1.7.0.
30. O'Hagan Anthony. SHELF: the Sheffield Elicitation Framework 2019. <http://www.tonyohagan.co.uk/shelf/ecourse.html> (accessed on 1 January 2021).
31. O'Hagan Anthony. Probabilistic Judgements for Expert Elicitation (e-learning course) 2018. Available at <http://www.tonyohagan.co.uk/shelf/ecourse.html>.
32. Oakley Jeremy E., O'Hagan Anthony. *SHELF: the Sheffield Elicitation Framework (version 4)*. School of Mathematics and Statistics, University of Sheffield, UK 2019. Available at <http://tonyohagan.co.uk/shelf>.
33. Neuenschwander Beat, Capkun-Niggli Gorana, Branson Michael, Spiegelhalter David J. Summarizing historical information on controls in clinical trials *Clinical Trials*. 2010;7:5-18. PMID: 20156954.
34. Gokhale D. V., Press S. James. Assessment of a Prior Distribution for the Correlation Coefficient in a Bivariate Normal Distribution *Journal of the Royal Statistical Society: Series A (General)*. 1982;145:237-249.
35. Trivedi Pravin K, Zimmer David M. *Copula Modeling: An Introduction for Practitioners*. Hanover: now Publishers Inc. 2007.
36. Erpenbeck Veit J., Popov Todor A., Miller David, et al. The oral CRTh2 antagonist QAW039 (fevipiprant): A phase II study in uncontrolled allergic asthma *Pulmonary Pharmacology & Therapeutics*. 2016;39:54–63.
37. Brightling Christopher E, Bleeker Eugene R, Erpenbeck Veit J, et al. LUSTER-1 and -2: Two randomized controlled trials of the prostaglandin D2 receptor 2 antagonist, fevipiprant, in asthma *Clinical Investigation*. 2019;9:55–63.

38. Petsky H L, Cates C J, Lasserson T J, et al. A systematic review and meta-analysis: tailoring asthma treatment on eosinophilic markers (exhaled nitric oxide or sputum eosinophils) *Thorax*. 2010;67:199–208.
39. Green Ruth H, Brightling Christopher E, McKenna Susan, et al. Asthma exacerbations and sputum eosinophil counts: a randomised controlled trial *The Lancet*. 2002;360:1715–1721.
40. Chlumský J, Striz I, Terl M, Vondracek J. Strategy Aimed at Reduction of Sputum Eosinophils Decreases Exacerbation Rate in Patients with Asthma *Journal of International Medical Research*. 2006;34:129–139.
41. Jayaram L.. Determining asthma treatment by monitoring sputum cell counts: effect on exacerbations *European Respiratory Journal*. 2006;27:483–494.
42. Nair Parameswaran, Pizzichini Marcia M.M., Kjarsgaard Melanie, et al. Mepolizumab for Prednisone-Dependent Asthma with Sputum Eosinophilia *New England Journal of Medicine*. 2009;360:985–993.
43. Fleming Louise, Wilson Nicola, Regamey Nicolas, Bush Andrew. Use of sputum eosinophil counts to guide management in children with severe asthma *Thorax*. 2011;67:193–198.
44. Castro Mario, Mathur Sameer, Hargreave Frederick, et al. Reslizumab for Poorly Controlled, Eosinophilic Asthma *American Journal of Respiratory and Critical Care Medicine*. 2011;184:1125–1132.
45. Pavord Ian D, Korn Stephanie, Howarth Peter, et al. Mepolizumab for severe eosinophilic asthma (DREAM): a multicentre, double-blind, placebo-controlled trial *The Lancet*. 2012;380:651–659.
46. Laviolette Michel, Gossage David L., Gauvreau Gail, et al. Effects of benralizumab on airway eosinophils in asthmatic patients with sputum eosinophilia *Journal of Allergy and Clinical Immunology*. 2013;132:1086–1096.e5.
47. Wenzel Sally, Ford Linda, Pearlman David, et al. Dupilumab in Persistent Asthma with Elevated Eosinophil Levels *New England Journal of Medicine*. 2013;368:2455–2466.
48. Ortega Hector G., Liu Mark C., Pavord Ian D., et al. Mepolizumab Treatment in Patients with Severe Eosinophilic Asthma *New England Journal of Medicine*. 2014;371:1198–1207.
49. Gauvreau Gail M., O'Byrne Paul M., Boulet Louis-Philippe, et al. Effects of an Anti-TSLP Antibody on Allergen-Induced Asthmatic Responses *New England Journal of Medicine*. 2014;370:2102–2110.

50. Castro Mario, Zangrilli James, Wechsler Michael E, et al. Reslizumab for inadequately controlled asthma with elevated blood eosinophil counts: results from two multicentre, parallel, double-blind, randomised, placebo-controlled, phase 3 trials *The Lancet Respiratory Medicine*. 2015;3:355–366.
51. Bleecker Eugene R, FitzGerald J Mark, Chanez Pascal, et al. Efficacy and safety of benralizumab for patients with severe asthma uncontrolled with high-dosage inhaled corticosteroids and long-acting  $\beta_2$ -agonists (SIROCCO): a randomised, multicentre, placebo-controlled phase 3 trial *The Lancet*. 2016;388:2115–2127.
52. FitzGerald J Mark, Bleecker Eugene R, Nair Parameswaran, et al. Benralizumab, an anti-interleukin-5 receptor  $\alpha$  monoclonal antibody, as add-on treatment for patients with severe, uncontrolled, eosinophilic asthma (CALIMA): a randomised, double-blind, placebo-controlled phase 3 trial *The Lancet*. 2016;388:2128–2141.
53. Corren Jonathan, Parnes Jane R., Wang Liangwei, et al. Tezepelumab in Adults with Uncontrolled Asthma *New England Journal of Medicine*. 2017;377:936–946.
54. Panettieri Reynold A, Sjöbring Ulf, Péterffy AnnaMaria, et al. Tralokinumab for severe, uncontrolled asthma (STRATOS 1 and STRATOS 2): two randomised, double-blind, placebo-controlled, phase 3 clinical trials *The Lancet Respiratory Medicine*. 2018;6:511–525.
55. Russell Richard J, Chachi Latifa, FitzGerald J Mark, et al. Effect of tralokinumab, an interleukin-13 neutralising monoclonal antibody, on eosinophilic airway inflammation in uncontrolled moderate-to-severe asthma (MESOS): a multicentre, double-blind, randomised, placebo-controlled phase 2 trial *The Lancet Respiratory Medicine*. 2018;6:499–510.
56. Castro Mario, Corren Jonathan, Pavord Ian D., et al. Dupilumab Efficacy and Safety in Moderate-to-Severe Uncontrolled Asthma *New England Journal of Medicine*. 2018;378:2486–2496.
57. Holzhauer Björn, Wang Craig, Schmidli Heinz. Evidence synthesis from aggregate recurrent event data for clinical trial design and analysis *Statistics in Medicine*. 2017;37:867–882.
58. Zondervan-Zwijnenburg Mariëlle, Schoot-Hubeek Wenneke, Lek Kimberley, Hoijsink Herbert, Schoot Rens. Application and Evaluation of an Expert Judgment Elicitation Procedure for Correlations *Frontiers in Psychology*. 2017;8:90.