



What if we Find Nothing? Bayesian Analysis of the Statistical Information of Null Results in Future Exoplanet Habitability and Biosignature Surveys

Daniel Angerhausen^{1,2,3} , Amedeo Balbi⁴ , Andjelka B. Kovačević⁵ , Emily O. Garvin¹ , and Sascha P. Quanz^{1,6} 

¹ETH Zurich, Institute for Particle Physics & Astrophysics, Wolfgang-Pauli-Str. 27, 8093 Zurich, Switzerland; dangerhau@ethz.ch

²Blue Marble Space Institute of Science, Seattle, WA, USA

³SETI Institute, 189 N. Bernardo Ave, Mountain View, CA 94043, USA

⁴Dipartimento di Fisica, Università di Roma “Tor Vergata,” Italy

⁵University of Belgrade-Faculty of Mathematics, Department of Astronomy, Studentski trg 16, 11000 Belgrade, Serbia

⁶ETH Zurich, Department of Earth and Planetary Sciences, Sonneggstrasse 5, 8092 Zurich, Switzerland

Received 2024 December 19; revised 2025 February 19; accepted 2025 February 22; published 2025 April 7

Abstract

Future telescopes will survey temperate, terrestrial exoplanets to estimate the frequency of habitable (η_{Hab}) or inhabited (η_{Life}) planets. This study aims to determine the minimum number of planets (N) required to draw statistically significant conclusions, particularly in the case of a null result (i.e., no detections). Using a Bayesian framework, we analyzed surveys of up to $N = 100$ planets to infer the frequency of a binary observable feature (η_{obs}) after null results. Posterior best fits and upper limits were derived for various survey sizes and compared with predicted yields from missions like the Large Interferometer for Exoplanets (LIFE) and the Habitable Worlds Observatory (HWO). Our findings indicate that $N = 20\text{--}50$ “perfect” observations (100% confidence in detecting or excluding the feature) yield conclusions relatively independent of priors. To achieve 99.9% upper limits of $\eta_{\text{obs}} \leq 0.2/0.1$, approximately $N \simeq 40/80$ observations are needed. For “imperfect” observations, uncertainties in interpretation and sample biases become limiting factors. We show that LIFE and HWO aim for sufficiently large survey sizes to provide statistically meaningful estimates of habitable environments and life prevalence under these assumptions. However, robust conclusions require careful sample selection and high-confidence detection or exclusion of features in each observation.

Unified Astronomy Thesaurus concepts: [Exoplanets \(498\)](#); [Surveys \(1671\)](#); [Bayesian statistics \(1900\)](#)

1. Introduction

In the past two decades, the discovery of thousands of exoplanets has led to increasingly accurate estimates of the occurrence of rocky worlds in the habitable zone of their stars. At the same time, the focus of observations has been gradually shifting from detection to characterization, with a number of current and future instruments, and missions devoted to refining the sample of potentially habitable nearby planets and to perform detailed spectroscopic study of their atmospheres. These include the Transiting Exoplanet Survey Satellite (TESS; G. R. Ricker et al. 2014), the CHAracterising ExOPlanets Satellite (W. Benz et al. 2021), PLANetary Transits and Oscillations of stars (H. Rauer et al. 2014, Atmospheric Remote-sensing Infrared Exoplanet Large-survey (G. Tinetti et al. 2021) space missions, the ground-based European-led Extremely Large Telescope (ELT; R. Gilmozzi & J. Spyromilio 2007), JWST (T. P. Greene et al. 2016), as well as future projects currently under study, such as the Habitable Worlds Observatory (HWO; National Academies of Sciences, Engineering, and Medicine 2021) and the Large Interferometer for Exoplanets (LIFE; S. P. Quanz et al. 2022b). The ambitious long-term goal of such observations is assessing the actual habitability of temperate, terrestrial exoplanets in our galactic neighborhood, and possibly detecting spectroscopic signatures of global biospheres (S. Seager 2014; J. L. Grenfell 2017; Y. Fujii et al. 2018; E. W. Schwieterman et al. 2018).

One of the great accomplishments of exoplanets science in the past years was the derivation of first estimates of η_{earth} —the mean number of earth-like rocky planets per star—using the large number of exoplanet detections made by the Kepler (W. J. Borucki et al. 2010) and TESS satellites, as well as ground-based transit and radial velocity surveys (see, e.g., R. K. Kopparapu et al. 2018; S. Bryson et al. 2021; G. J. Bergsten et al. 2022). In analogy to this, one of the prospective outputs of the surveys discussed here, will be to derive estimates of the frequency of habitable η_{Hab} or even inhabited planets η_{Life} .

First empirical estimates for η_{Hab} or η_{Life} will certainly have an enormous impact. However, it is complicated to observationally constrain these quantities by remote detections of limited samples alone (see below and in the discussion in Section 4). These variables are among the major remaining unknowns essential for evaluating the prevalence of life in the universe. For example, they play a pivotal role in the renowned Drake equation and its modifications (F. Drake 1965; S. Seager 2018). Furthermore, these estimates have a direct impact on gauging the potential future lifespan of humanity, influencing assessments of possible bottlenecks, often referred to as the “Great Filter,” in the evolution of life (A. Chopra & C. H. Lineweaver 2016; J. Haqq-Misra et al. 2020). For instance, the surveys discussed here could show that even on a planet with suitable conditions, abiogenesis might still be an extremely rare event. If these future observations on the other hand reveal that simple life is abundant, but complex or intelligent life is absent, it would suggest that subsequent evolutionary steps (e.g., multicellularity, intelligence, or

technological advancement) are suppressed by the “Great Filter.”

Unfortunately, our current knowledge is restricted only to the emergence of life on Earth, making it exceedingly challenging to draw meaningful inferences about the prevalence of life elsewhere. This limitation is compounded by the presence of selection biases, which render very problematic to make conclusive assessments (A. Balbi & M. Lingam 2023). Nevertheless, some limited insights can be gleaned from the relatively rapid appearance of life on our planet (D. S. Spiegel & E. L. Turner 2012; D. Kipping 2020). In this context, it is important to highlight additional relevant insights and investigations that contribute to our understanding of planetary systems, the factors influencing habitability, and the broader search for extraterrestrial life. Examples are planet formation models that predict the occurrence of habitable planets (e.g., R. Burn et al. 2021; A. Emsenhuber et al. 2021; T. Kimura & M. Ikoma 2022), the “red sky paradox,” which questions why Earth orbits a G-type star rather than a later-type star (D. Kipping 2021), studies of the prevalence of planetary conditions that may support abiogenesis (e.g., S. A. Benner 2014; S. Ranjan & D. D. Sasselov 2016; A. Longo & B. Damer 2020; D. D. Sasselov et al. 2020; D. Deamer 2022) and/or habitable conditions (e.g., E. S. Kite & E. B. Ford 2018; J. Krissansen-Totton et al. 2022; M. Schlecker et al. 2024), or formal birth–death relationships (D. Kipping & G. Lewis 2025).

Building on previous studies such as A. Balbi & C. Grimaldi (2020) and (S. P. Quanz et al. 2022a), which analyzed the potential impact of future biosignature surveys on our understanding of the frequency of inhabited worlds in our galaxy, this work offers a novel perspective.

We aim to estimate the minimum number of exoplanets required for future observations to yield statistically significant conclusions regarding η_{Hab} and η_{Life} . Our focus is specifically on the clearly defined scenario of a null result, i.e., where no surveyed planet shows signs of habitability or life. Without diverting into a detailed discussion, individual observations of the survey that would actually indicate a habitable planet could be the detection of ocean glint (T. D. Robinson et al. 2010; J. Lustig-Yaeger et al. 2018), strong heat redistribution (Y. Fujii et al. 2025, in preparation; M. Braam et al. 2025, in preparation) on a tidally locked planet or certain amounts of water vapor in the atmosphere (S. Rugheimer et al. 2025, in preparation) that point at the presence of surface water. For the individual (unambiguous) detection of life on these planets it could be observations of (capstone) biosignatures (e.g., D. Angerhausen et al. 2024) or even technosignatures (e.g., E. W. Schwieterman et al. 2024). Initially, we adopt an assumption of “perfect knowledge” in our individual observations, implying absolute certainty in the absence of detected features, thus eliminating the possibility of false negatives or positives. This baseline helps us understand the upper limits of potential results of surveys, when being absolutely certain that each observation accurately reflects the true state of the planet. Similar methodological issues have been discussed in different contexts such as general statistics, particle physics, and high-energy astrophysics (see e.g., L. D. Brown et al. 2001; G. Cowan 2007; V. L. Kashyap et al. 2010).

Of course, this is a highly idealized scenario, as remarked in several earlier studies. For instance, A. R. Truitt et al. (2020), D. C. Catling et al. (2018), and S. I. Walker et al. (2018) presented (Bayesian) frameworks to quantify the confidence in

individual observations, while H. Smith & C. Mathis (2023), S. Foote et al. (2023), and T. Fisher et al. (2023) discussed the problems with false positives/negatives, underscoring that “high confidence in life detection claims require either (1) a strong prior hypothesis about the existence of life in a particular alien environment, or conversely, (2) signatures of life that are not susceptible to false positives.” This highlights the importance of prior probabilities and robust likelihood functions. Furthermore, P. Vickers et al. (2023) addressed the problem of unconceived alternative explanations of otherwise agreed upon biosignatures by emphasizing the need to account for unknown hypotheses in the Bayesian model. J. Green et al. (2021) and C. Gillen et al. (2023) discussed solutions to better quantify and communicate the confidence in individual (non) detections using Bayesian credible intervals and posterior distributions. We attempt to take these intricacies into account in Section 2.4 of our analysis, where we model the observational uncertainty by introducing a probability of obtaining false-negative results or prior knowledge about a biased sample.

2. Method

The main questions we want to address in our study are these: assuming we observed a number N of planets in a given survey and can exclude a certain feature (i.e., habitability or the presence of life) for all of them, how well can we quantify the universal fraction η_{obs} of planets having this feature (e.g., η_{Hab} or η_{Life})? Also, how does it impact our ability to derive these fractions if we are not 100% confident on each single observation in our survey of this feature?

We tackle this problem by simulating a sample of N individual observations, with N varying between 1 and 100. This allows us to assess the significance of the results, by asking for which number N of observed planets we can cross certain thresholds in our derived knowledge about the absolute fraction η_{obs} having the desired feature.

We first adopt a Bayesian approach and use a set of assigned prior beliefs for η_{obs} (“optimistic,” “neutral,” or “pessimistic”; see Section 2.2) to compute the posterior belief distribution while conditioning on the findings of null results. The results from the posterior distribution are obtained over N observations using Bayesian updating procedures. We use the posterior distribution to derive the credible intervals for η_{obs} as a function of N , and identify the minimal sample size requirements for a given level of assigned posterior belief. We then also compare to a frequentist approach to the same set of questions (see Section 2.6).

2.1. The Beta-binomial Model

In this study we are dealing with observations asking a “yes or no” question in each individual element of our survey. Therefore we adopt a beta distribution to represent our preliminary understanding/prior belief $f(\eta_{\text{obs}})$ of the fraction of a binomially distributed observational parameter η_{obs} (for example η_{Hab} or η_{Life}) $\in [0, 1]$. The beta distribution is a suitable choice for modeling our prior distribution of η_{obs} for several reasons. First, it is defined on the interval $[0, 1]$, aligning with the range of probabilities. Additionally, it serves as the conjugate prior for the likelihood which is defined as a binomial distribution, a primary choice for modeling a sequence of independent observations, each having a binary




Figure 1. Comparison of the various priors used in this study. While the Bayes–Laplace (“flat”) prior might seem “uninformative” when plotted in linear space, the logarithmic scale shows how it “optimistically” puts a lot of weight in values of $\eta_{\text{obs}} \sim 1$. This might not be the best choice for astrobiological questions, if one assumes that processes are either extremely rare or very frequent. Therefore we also chose more “neutral” priors such as the Jeffreys or Kerman ones, and also one extremely “pessimistic” one, giving more weight to $\eta_{\text{obs}} \approx 0$.

outcome. The main benefit of a conjugate prior is that it gives a closed-form expression for the posterior, significantly reducing the numerical burden. In general conjugate priors give a very intuitive insight into how a likelihood function updates a prior distribution. Lastly, the beta distribution has the flexibility to capture a diverse spectrum of prior beliefs regarding the probability of success. We then update our beliefs about η_{obs} based on the results of our survey of exoplanets in our sample for each $N \in [1, 100]$.

The general definition of the beta distribution is:

$$f(\eta_{\text{obs}}|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \eta_{\text{obs}}^{\alpha-1} (1 - \eta_{\text{obs}})^{\beta-1} \quad (1)$$

where Γ denotes the gamma function. In the following sections we explain how α and β are defined in the particular context of our analysis. From the practical perspective, we choose the conjugate beta-binomial model for its computational simplicity and ease of interpretation. It makes the calculations of Bayesian posteriors extremely straightforward (e.g., A. Johnson et al. 2022), described in Section 2.3.

2.2. Defining Priors

Prior functions are used in Bayesian statistics to specify our belief in the possible distribution of a variable before performing the experiment. The prior can be more or less informative, depending on the extent of previous knowledge it conveys regarding the variable. In situations where very limited or no a priori information is available, an “uninformative” prior is typically employed. It is essential to recognize that even in such cases, a small degree of objective information is factored into the analysis, such as the range within which the variable is defined.

As we mentioned earlier, the beta function can represent varying prior distributions depending on the choice of the parameters α and β . In our analysis, we adopted the following five choices of priors (see Figure 1):

1. *Bayes–Laplace (“flat”) prior.* This prior assumes that all values of the variable are equally likely. It is given by the beta distribution with $\alpha_{\text{flat}} = 1$ and $\beta_{\text{flat}} = 1$. Although this prior is sometimes considered “uninformative,” it is in fact rather optimistic in our context, as it gives a large relative weight to values of η_{obs} close to unity.
2. *Jeffreys prior.* This prior is a more appropriate choice for the “uninformative” case, as it gives a similar weight to the extreme cases of $\eta_{\text{obs}} \approx 0$ and $\eta_{\text{obs}} \approx 1$. It is derived from the Fisher information matrix, and it is given by the beta distribution with $\alpha_{\text{jeff}} = 0.5$ and $\beta_{\text{jeff}} = 0.5$.
3. *Kerman prior (J. Kerman 2011).* This prior falls roughly in between the Jeffreys prior and the pessimistic prior and can be considered a “neutral” prior, because it is constructed so that the posterior distributions contain approximately 50% probability that the true value is either smaller or larger than the maximum likelihood estimate. It is given by the beta distribution with $\alpha_{\text{ker}} = 0.333$ and $\beta_{\text{ker}} = 0.333$.
4. *Solar system “biased” optimist.* This prior is informative. It is given by the beta distribution with $\alpha_{\text{opt}} = 2$ and $\beta_{\text{opt}} = 3$. This prior would for example be derived from a posterior after using a flat prior ($\alpha_{\text{flat}} = \beta_{\text{flat}} = 1$) and the solar system’s sample of terrestrial, temperate (“optimistic habitable zone”) planets as our baseline sample. In that case we would add 1 positive observation (Earth) to α and 2 negative observations (Mars, Venus) to β when deriving the general “fraction of planets in the optimistic HZ that are habitable/inhabited.” In that sense it would represent an interpretation suffering from selection bias (we can only exist as observers on a living planet) leading to a presumably too optimistic prior. While this prior is still relatively low in information/belief by only slightly preferring solutions around ~ 0.3 , it tends to produce rather “optimistic” posteriors (i.e., still believing in a relatively high fraction of living planets) in comparison to other more “pessimistic” priors (i.e., preferring solution with very low frequencies) when given the same set of

updated information through additional nondetection observations (see also, in Appendix B, Figure 6.)

5. *Extreme pessimistic prior.* With weights mainly close to 0 and 1 this prior is also very informative in the sense that it assumes that the given feature will either be found for all or none of the observations. This and similar priors (e.g., the classic Haldane prior)⁷ are discussed and used widely in astrobiological reasoning, also, for example, in the context of the “fine-tuning problem” (D. Kipping & G. Lewis 2025). Once presented with only a few negative results, the posteriors belief is that the probability of success is very small, by giving a large relative weight to values of η_{obs} close to 0. It is given by the beta distribution with $\alpha_{\text{ext}} = 0.005$ and $\beta_{\text{ext}} = 0.005$.

Figure 6 in Appendix B exemplifies and illustrates why the flat prior can be considered optimistic: for one positive observation in a hundred tests, the resulting posterior is strongly peaked around 1%, being quite confident that this is the actual fraction; conversely, for the extreme pessimistic prior, even for one in a hundred detection the posterior still assigns strong credence to very low fractions, $<10^{-4}$). In this context, the Jeffreys and Kerman prior can be considered “neutral,” i.e., not too optimistic or pessimistic.

2.3. Simulation of Observational Scenarios

We use Bayesian inference to calculate η_{obs} , the fraction of planets possessing some binary observable feature, given a prior belief about this fraction and update our posteriors with increasing number N of observations. Our prior belief is modeled by the above mentioned beta distributions with shape parameters α_{prior} and β_{prior} .

The posterior distribution for the fraction of planets having the desired feature, after n_{pos} positive cases have been found in N observations, is then calculated as follows:

$$\begin{aligned} f(\eta_{\text{obs}} | \alpha_{\text{prior}}, \beta_{\text{prior}}, N, n_{\text{pos}}) \\ = \frac{\Gamma(\alpha_N + \beta_N)}{\Gamma(\alpha_N)\Gamma(\beta_N)} \eta_{\text{obs}}^{\alpha_N-1} (1 - \eta_{\text{obs}})^{\beta_N-1} \end{aligned} \quad (2)$$

with $\alpha_N = \alpha_{\text{prior}} + n_{\text{pos}}$ and $\beta_N = \beta_{\text{prior}} + N - n_{\text{pos}}$.

In other words, to update our interpretation of the survey results we would add the number of positive observations to α and the number of nondetections to β to derive the posteriors at that stage of the survey.

In the first part of our simulations, we assume a survey of “perfect” nondetection, i.e., $n_{\text{pos}} = 0$ with no ambiguity. We then update the prior distribution varying the survey size N between 1 and 100, leading to posteriors in the form of beta $(\alpha, \beta + N)$.

2.4. Modeling Imperfect Individual Observations

In reality, each individual observation (asking does my planet have the feature or not?) will have a degree of uncertainty. In this section we want to quantify the extent of this uncertainty and its potential impact on our results and statistical power, particularly under worst-case scenarios. But what does it actually mean if someone is “80% sure they did not find water” or concludes that they “can exclude the

presence of life with 3σ ?” Can we even find parameterizations to quantify this?

Bayesian inference allows us to update our beliefs about a parameter in light of new data. When the beta distribution is used as a conjugate prior for a binomial likelihood, the posterior distribution remains a beta distribution, enabling straightforward updating as new data is observed. Typically, this Bayesian update involves adjusting the parameters of the beta distribution based on the number of successes (n_{pos}) and the total number of observations (N). However, in the presence of uncertainties—such as sample misclassification or potential false negatives—the effective amount of information provided by the data may be reduced or altered. To account for this, we adjust the parameters of the beta distribution to reflect the true information content of the data, rather than treating all observations as equally reliable. This adjustment is based on the concept of effective sample size (see, e.g., P.-J. F. Hulson et al. 2011), which quantifies the information that a sample provides about a population parameter, accounting for dependencies among observations, variability, and other sources of uncertainty. In our analysis, we adjust for sample uncertainty (U_s) and interpretation uncertainty (U_i). Sample uncertainty (U_s) reflects the belief that some observed planets may not belong to the target sample. For example, if a mini-Neptune is mistakenly included in a sample of rocky planets, it should not influence the estimation of (η_{obs}), the fraction of rocky planets possessing the desired feature. To account for this, each observation is weighted by the factor $(1 - U_s)$, where (U_s) represents the level of sample uncertainty. The posterior distribution is then updated as follows:

$$\eta_{\text{obs}} | \text{data} \sim \text{Beta}(\alpha + n_{\text{pos}}(1 - U_s), \beta + N(1 - U_s) - n_{\text{pos}}(1 - U_s)). \quad (3)$$

In cases where no positive detections are observed ($n_{\text{pos}} = 0$), the posterior distribution is given by:

$$\eta_{\text{obs}} | \text{data} \sim \text{Beta}(\alpha, \beta + N(1 - U_s)). \quad (4)$$

This is analogous to using a weighted likelihood, where the effective sample size is reduced to account for the uncertainty in the relevance of the data. By doing so, we ensure that the posterior distribution accurately reflects the confidence we have in the sample’s composition. Interpretation uncertainty (U_i), on the other hand, addresses the possibility that some planets are incorrectly labeled as not having the desired feature, leading to false negatives. To address this, we adjust the posterior by adding a fraction (NU_i) of the total observations as “hidden” positives to the (α) parameter, while simultaneously reducing the number of true negatives added to β by (NU_i). The posterior distribution is then updated as follows:

$$\eta_{\text{obs}} | \text{Data} \sim \text{Beta}(\alpha + NU_i + n_{\text{pos}}, \beta + N(1 - U_i)). \quad (5)$$

For cases with no positive detections ($n_{\text{pos}} = 0$), the posterior distribution becomes:

$$\eta_{\text{obs}} | \text{Data} \sim \text{Beta}(\alpha + NU_i, \beta + N(1 - U_i)). \quad (6)$$

This approach is akin to a robust Bayesian analysis, where the model accounts for potential errors in data classification. By incorporating “hidden” positives, we adjust for the possibility of false negatives, leading to a posterior distribution that more accurately reflects the underlying parameter (η_{obs}). By

⁷ We do not use the Haldane prior ($\alpha = \beta = 0$) because it yields improper posterior distributions for the extreme outcomes such as N or 0 positive observations in N trials.

adjusting for sample misclassification and data interpretation errors, we reduce the likelihood of bias in the posterior estimates.

In a (maybe more intuitive) sense these two approaches represent two extreme possibilities: for the U_s an uncertainty of 0.1 would mean that each (negative) observation only counts with a weigh of 0.9. For the U_i scenario; however, an uncertainty of 0.1 would mean that each (negative) observation only counts with a weigh of 0.9 and also add a 0.1 fraction of a positive observation to the interpretation.

As an illustration, let us consider a scenario where we conducted 50 observations without observing the desired feature. If we are confident in the outcome of each observation at the 80% level, the most pessimistic assumption would be that we were wrong 20% of the times (i.e., that U_i was 0.2), corresponding, on average, to 10 false negatives in our sample. The most optimistic interpretation would be that in these 20% of cases we could just not say anything (i.e., as if the planet wasn't even from our sample, $U_s = 0.2$). Practical examples for these two cases could be: (1) a priori knowledge (e.g., from completeness studies of the survey) of a certain rate of wrongly labeled planets in the sample; or (2) results from combined retrievals and (biological) interpretation that derive a certain rate of mislabeled planets that actually inhabit the desired feature.

We model both with three different levels of U_s (50%, 33%, and 10%) and U_i (20%, 10%, and 5%) .

2.5. Calculating Limits

We then calculate the upper and lower limits of the posterior distribution at each “step” of the survey, i.e., after each of the N observations, by finding the values $f_{\text{obs,low}}$ and $f_{\text{obs,up}}$ such that:

$$\int_0^{\eta_{\text{obs,low}}} f(\eta_{\text{obs}} | \alpha_{\text{prior}}, \beta_{\text{prior}}, N, n_{\text{pos}}) d\eta_{\text{obs}} = 0.001, \quad (7)$$

and

$$\int_0^{\eta_{\text{obs,up}}} f(\eta_{\text{obs}} | \alpha_{\text{prior}}, \beta_{\text{prior}}, N, n_{\text{pos}}) d\eta_{\text{obs}} = 0.999. \quad (8)$$

2.6. Comparison to Frequentist Benchmark

We compare the results of the Bayesian approach with those of a frequentist approach to provide a comprehensive analysis and ensure consistency across methods. It also allows us to show that our general conclusions are method independent. The key distinction of the frequentist approach is that it does not incorporate prior beliefs: probabilities are solely derived from the observed event frequencies. Hence, let n_{pos} represent the number of successes resulting from N independent trials with unknown success probability η , where the random variable n_{pos} follows a binomial distribution (as a repetition of Bernoulli trials):

$$n_{\text{pos}} \sim \text{Bin}(N, \eta) \quad (9)$$

with the following likelihood:

$$L(\eta | N, n_{\text{pos}}) = \binom{N}{n_{\text{pos}}} \eta^{n_{\text{pos}}} (1 - \eta)^{(N - n_{\text{pos}})}. \quad (10)$$

Due to the absence of observations of positive events ($x = 0$), the maximum likelihood estimator of η_{obs}

$$\hat{\eta}_{\text{mle}} = \frac{n_{\text{pos}}}{N} \quad (11)$$

will be 0. Thus, to overcome related computational shortcomings while building the confidence intervals, we employ an exact binomial confidence interval, with a lower limit at 0 and the upper limit calculated from the cumulative distribution function of the binomial based on a 99.9% one-sided confidence interval. We also compare it with a conservative 99.9% Wilson confidence interval based on the normal approximation, to ensure reliability of our results under larger sample sizes (for $N \geq 30$). Thus, the lower bound of the Wilson confidence interval is defined as $CI_{\text{wilson},0.999}$ falls at 0. The upper bound is defined as $CI_{\text{wilson},0.999}^* = \min(1, CI_{\text{wilson},0.999}^*)$, with

$$CI_{\text{wilson},0.999}^* = \frac{1}{1 + \frac{z_{\alpha}^2}{N}} \times \left(\hat{\eta}_{\text{mle}} + \frac{z_{\alpha}^2}{2N} \pm \frac{z_{\alpha}}{2N} \sqrt{4N\hat{\eta}_{\text{mle}}(1 - \hat{\eta}_{\text{mle}}) + z_{\alpha}^2} \right). \quad (12)$$

3. Results

3.1. Observations with “Perfect” Individual Nondetection

The results for surveys with “perfect”⁸ nondetection for each of the N individual observations are shown in Figure 2. There we show the posterior mean (“best fit”) from the Bayesian analysis described above for the fraction η_{obs} as a function of the number N of nondetections, as well as the 99.9% upper limits derived from the credible intervals of the posteriors, for the different prior beliefs. We chose 99.9% as our threshold value to guide the intuition (particularly in public outreach or policy making contexts) that “there is only a 1 in 1000 chance” that the *real* number is *not* below our derived upper limit. Also for an approximately normal dataset, this threshold corresponds to a value slightly beyond three standard deviations ($3\sigma \sim 99.7\%$). Around $N = 20$ all priors “agree” on a “best fit” lower than ~ 0.05 , while after ~ 40 observations they all “agree” on $\eta_{\text{obs}} < 0.15$ with 99.9% confidence.

In Figure 2 (right, dashed/dotted) we also show a comparison with the results one would get using the frequentist approaches to the problem. The confidence interval of the exact binomial aligns with the credible interval from the flat and uninformative prior, while the Wilson confidence interval aligns to the optimist credible interval for large sample regime. Although we plot those upper limits together with Bayesian credible intervals, it is important to note that confidence intervals and credible intervals encapsulate alternative meanings and probabilistic interpretations (S. Jin et al. 2017). Nevertheless, they are expected to converge in asymptotic behavior.

In summary, we would need more than ~ 40 or 80 individual, “perfect” observations, respectively, to exclude scenarios with $\eta_{\text{obs}} > 0.2$ and 0.1 with 99.9% belief (or confidence in the frequentist framework). A similarly large

⁸ For example, no uncertainties in our sample population or observation interpretation.

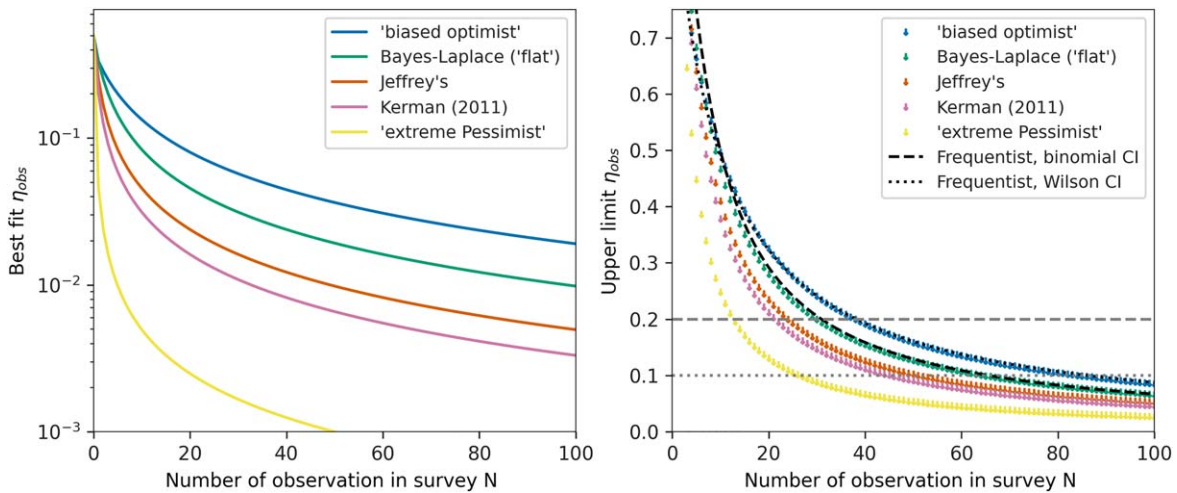


Figure 2. Final results for the case of a null-result survey of “perfect” observations. Left: best fit from Bayesian analysis for the fraction η_{obs} as a function of the number N of nondetections. The frequentist best fit corresponds to $\hat{\eta}_{\text{MLE}} = 0$, and hence can not be represented in this plot. Right: 99.9% upper limits derived from the credible intervals of the posteriors for the different prior beliefs. Here we also compare Bayesian to Frequentist uncertainty assessments, by including both an exact binomial confidence interval, as well as a Wilson confidence interval based on a normal approximation. In order to exclude scenarios with $\eta_{\text{obs}} > 0.2$ and 0.1 with 99% belief, we would need more than ~ 40 or 80 individual, “perfect” observations, respectively. As expected and discussed in Section 2.2, the “optimistic” priors produce posteriors with higher values, whereas the “pessimistic” priors are at the lower end of the spectrum.

number of observations is needed to reduce the impact of the chosen prior to be smaller than 5%–10% (see Figure 7) on the final derived values.

3.2. Modeling Observational Uncertainty in Individual Nondetection

Figure 3 shows the best fit (left) as well as the corresponding 99.9% upper limits (right) as a function of the number of observations N in our survey, for different sample (top) and interpretation (bottom) uncertainties using the Jeffreys prior. The figure shows how cases with non-perfect observations, assuming a certain interpretation uncertainty, always produce higher values for both the best fit and the upper limit. Ultimately, the derived values are limited by U_i . In other words, if we observe a series of nondetections and (cautiously) assume in our interpretation that all observations beyond our confidence levels are wrongly interpreted, we can never derive values lower than the U_i . For instance, if we perform 100 observations assuming a 20% rate of false interpretations, our best fit will never be smaller than 0.2, and the 99.9% upper limit will never be smaller than 0.4, no matter how much we increase N . This highlights the fact that it is only possible to derive meaningful fractions η_{obs} if the observable feature is unambiguous and can be confirmed or excluded with confidence levels higher than the desired assigned belief on the final result for η_{obs} . It also emphasizes the need to “ask the right questions,” i.e., to not conduct surveys for which our confidence in the interpretation of the individual observational is very low. As shown in Figure 3, an interpretation uncertainty of only 5% has a worse impact on our ability to constrain an upper limit than a sample uncertainty of 50%. This is why we did not model U_i larger than 20%.

For the cases of sample uncertainty U_s the problem is less pronounced: in that case we can actually compensate for the out-of-sample planets in our survey by observing a larger number of targets. In other words, even if we observe a series of nondetections and have to a priori belief that 50% of all observations do not actually fall into our defined sample, we

can just double our sample to $2N$ to derive the same upper limits as an unbiased survey with a perfect sample.

3.3. Application to Future Survey Missions

In this section we compare our previous findings with real life scenarios for future mission concepts, namely LIFE and HWO. For completeness we want to also mention adaptive optics supported high-dispersion spectroscopy that will be possible with the Planetary Camera and Spectrograph for the Extremely Large Telescope (ELT/PCS; M. Kasper et al. 2021). A hypothetical 10 yr survey—assuming best case scenarios—with ELT/PCS could detect modern Earth-like oxygen levels in the atmospheres of up to 19 Earth-like exoplanet candidates (EECs) around bright, nearby M stars within 20 pc (K. K. Hardegree-Ullman et al. 2025).

3.3.1. The Large Interferometer for Exoplanets (LIFE)

LIFE is a proposed space-based mid-infrared nulling interferometer (S. P. Quanz et al. 2022b). This mission will be designed with the capability to investigate the atmospheric properties of a large sample of (primarily) terrestrial, temperate exoplanets. S. P. Quanz et al. (2022b), F. A. Dannert et al. (2022), and J. Kammerer et al. (2022) discussed the detection yield of a LIFE-like mission. J. Kammerer et al. (2022) considered a 2 m aperture at 5% throughput baseline scenario for LIFE, which would detect (and eventually also characterize) $N_{2\text{ m}/\text{EEC}} = 6_{-3}^{+7}$ EECs around FKG stars. EEC here is defined as a planet with instellation flux range as received in the conservative HZ and an upper radius limit of $1.5 R_{\oplus}$ but a lower limit of $0.8 a_p^{-0.5} R_{\oplus}$, where a_p is the planet semimajor axis in units of astronomical unit (C. C. Stark et al. 2019). Using an extended scenario with a 3.5 m aperture setup, J. Kammerer et al. (2022) derived $N_{3.5\text{ m}/\text{EEC}} = 16_{-9}^{+16}$.

For the “LIFE baseline” $N_{2\text{ m}/\text{EEC}} = 6$ case, with clear (i.e., 100% certain) nondetections of the observable feature the presented framework (using the Jeffreys prior) would derive: $\eta_{\text{obs}} = 0.071$ and a 99.9% upper limit of $\eta_{\text{obs}} < 0.57$. This means that in case of an unsuccessful search for “life as we

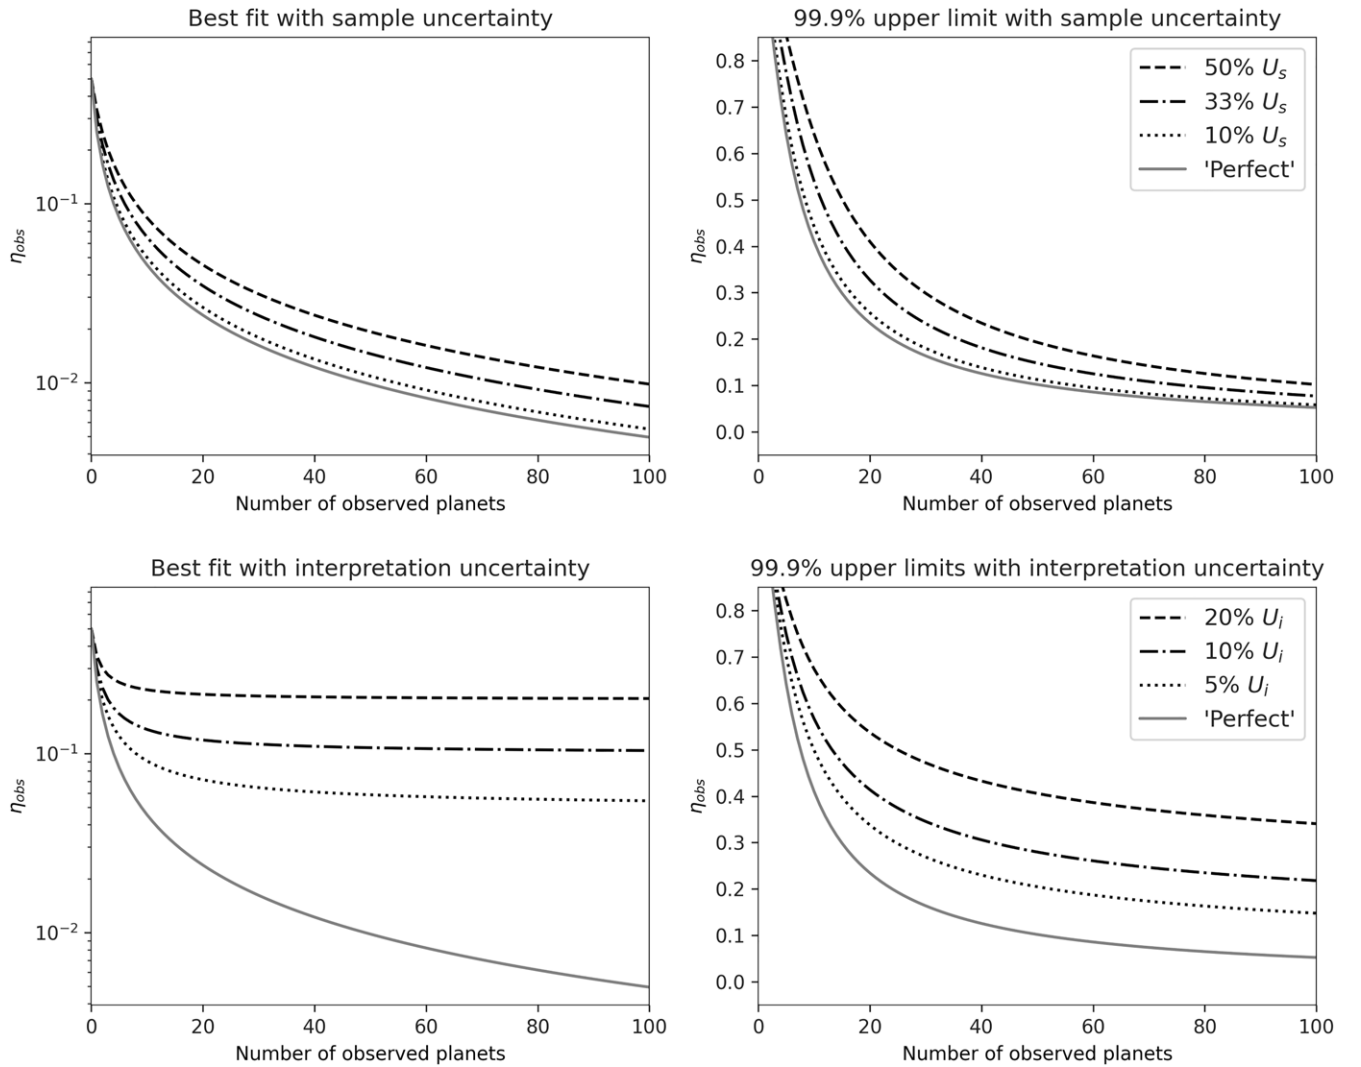


Figure 3. Testing different levels of “imperfect” individual observations, modeled as a sample (top) and interpretation uncertainty (bottom). Showing the difference between assuming “perfect” individual observations (solid line) and more realistic results with various levels of sample and interpretation uncertainty. For clarity we show only the calculations using the Jeffrey’s prior, as it seems to represent a good compromise between all the priors in the first general analysis (see, e.g., Figure 2).

know it on an Earth twin” in this sample of six planets we would (only) be able to constrain η_{Life} for FGK stars to be smaller than 57% with 99.9% belief.

For the second discussed “LIFE extended” $N_{3.5 \text{ m}/\text{EEC}} = 16$ case, the framework presented here (again, using the Jeffrey’s prior) would derive $\eta_{\text{obs}} = 0.029$ and a 99.9% upper limit of $\eta_{\text{obs}} < 0.28$. This means that in case of an unsuccessful search for habitable planets in this sample of 16 planets, LIFE would be able to constrain η_{Hab} for FGK stars to be smaller than 28% with 99.9% belief.

In Appendix A we compare different parameter spaces in radius and insulation also discuss LIFE’s capability of detecting and characterizing terrestrial planets in the optimistic HZ (a wider criterion than EEC) of FGK stars.

3.3.2. The Habitable Worlds Observatory (HWO)

NASA’s HWO⁹ is a flagship-class mission proposed by the US Astro2020 decadal survey (National Academies of

Sciences, Engineering, and Medicine 2021) that will utilize a combination of spectroscopic and imaging techniques to explore the atmospheres and surfaces of temperate and terrestrial exoplanets. For HWO we used numbers from R. Morgan et al. (2022)¹⁰: using a 6 m coronagraph scenario and constraining to “exo-Earths”¹¹ around “Sun-like” stars, they analyzed the yield of planets for which HWO could detect (or exclude) the presence of water or oxygen in the atmosphere based on a detection metric introduced there.

They show that they can measure a water observable metric for $N_{\text{H}_2\text{O}} = 20_{-2}^{+18}$ and an oxygen observable metric for $N_{\text{O}_2} = 17_{-2}^{+11}$ planets in the HWO survey. For $N_{\text{H}_2\text{O}} = 17$, our framework would derive $\eta_{\text{obs}} = 0.028$ and a 99.9% upper limit of $\eta_{\text{obs}} < 0.26$; for $N_{\text{O}_2} = 20$ we get $\eta_{\text{obs}} = 0.024$ and $\eta_{\text{obs}} < 0.23$; as before, we are assuming that no positive

⁹ Based on the the Large UV/Optical/IR Surveyor (The LUVOIR Team 2019), and the Habitable Exoplanet Observatory (B. S. Gaudi et al. 2020).

¹⁰ Their upper and lower limit are derived by the extreme cases of (i) no knowledge about existing planets before the start of their search phase or (ii) perfect knowledge of all available target which allows them to skip the search phase.

¹¹ Defined there as planets with $R \in [0.8\sqrt{a}, 1.4] R_{\oplus}$ orbiting at distances a with an insulation equivalent to $[0.95, 1.67]$ AU of the Sun, almost identical to the LIFE EEC criterion (see Appendix A).

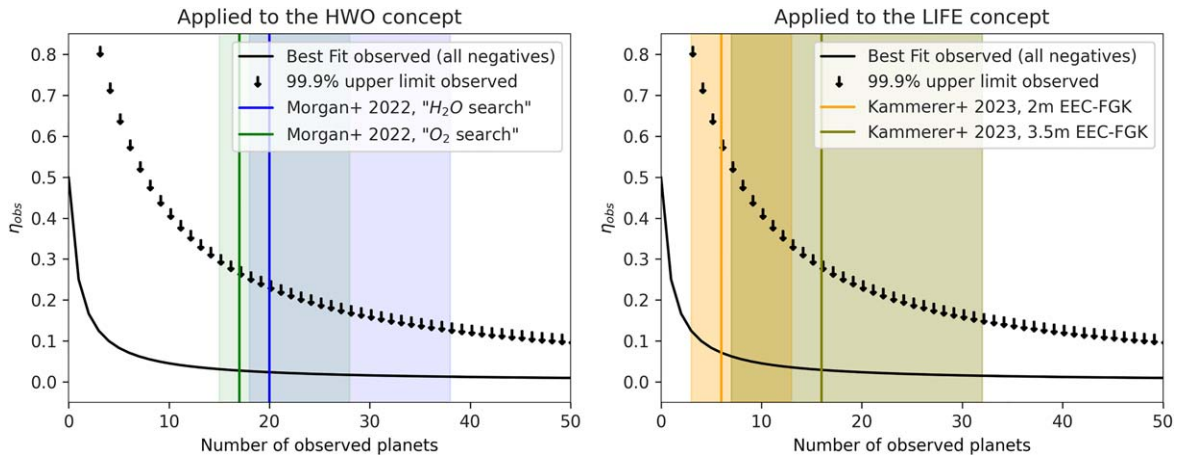


Figure 4. Comparison to potential survey sizes of the HWO mission and LIFE concept (using Jeffrey’s prior).

detection was made in the sample. This means that in case of an unsuccessful search for oxygen in the HWO sample of 20 planets, we would be able to constrain η_{O_2} for planets around FGK stars to be smaller than 23% with 99.9% belief, if none of the HWO observations will detect oxygen. Figure 4 shows the derived values for the whole range of possible N .

4. Discussion and Conclusions

In this work we show that future surveys of exoplanets aiming at characterizing several dozens of targets are large enough to derive meaningful statistical results for the frequency of observable features, even from a series of nondetections. Furthermore we show, that at these sample sizes, the results inferred from such surveys will be only weakly dependent on the choice of priors in a Bayesian interpretation of the result and comparable to conclusions made in a frequentist approach.

Nonetheless, it is essential to consider certain caveats and details in relation to these conclusions. First of all, we showed that our uncertainty in the individual interpretation of the observables has a strong impact on the belief in the final result, so that the uncertainties of individual observations must be very low in order to derive statistically significant frequencies. Another important point is to clearly define the sample (and understanding its limitations) as sample uncertainties also weaken the final conclusions. This means that, for example, asking which is the “fraction of planets that have life” will be an ill posed question for these purposes. Asking for “planets that have global biospheres” might be slightly better suited, since it at least implies the chance for remote detectability; however, it still relies on a strong consensus within the scientific community about which remote signatures indicate unambiguous detection of life (see D. C. Catling et al. 2018; V. Meadows et al. 2022; T. Fisher et al. 2023). From an observational standpoint, asking for something of the sort “planets within certain radii and temperature limits that show methane and oxygen simultaneously above certain thresholds” (see, e.g., S. D. Domagal-Goldman et al. 2014; V. S. Meadows et al. 2018) will be more feasible, in particular with regard to the ability to clearly say whether the feature was observed or not; framing the question in such a way, however, reduces the meaning and usefulness (in terms of generalizability) of the derived result. For example, how would we fit nonmethane/oxygen biosignatures in this framework, if detected? How generally applicable is η_{Hab} if potentially habitable sub-

Neptunes are excluded from the sample? Similar arguments could be applied to the sample on which the results are obtained: besides the general caveat that the planetary systems in our Solar neighborhood are not necessarily representative of all the system in other galactic environments, and that observations are not necessarily independent (e.g., assuming there is no panspermia on these distances, on multiple planets in the same system), biases in the sample itself (e.g., arbitrary cuts in stellar hosts), or the observation (e.g., it may be easier/clearer to detect habitability on certain types of planets) will limit the usefulness of the eventually derived fraction η_{obs} in more general contexts.

Despite the inherent challenges in estimating η_{Hab} or even inhabited η_{Life} with a small sample size, such as broader credible intervals and reduced precision, this still might offer several key useful aspects. First, preliminary insights gained from these estimates can serve as valuable guides for subsequent research directions in exoplanet studies, even if the data are limited. Establishing an initial baseline is crucial, as it provides a reference point for comparison with future, more comprehensive data sets. Moreover, tackling the complexities of small samples influences the development of sophisticated statistical methods. Understanding the limitations of small sample data is also instrumental in informing resource allocation for future observational efforts. These early estimates, while statistically limited, can significantly influence research policies and funding decisions. We are currently working on a follow up paper that generalizes the study conducted here to all sorts of false-negative and false-positive rates as well as underlying distributions of the observable feature to even further optimize future survey strategies.

Last but not least we want to remind the reader here that, even if this paper is about null results, a single positive detection would be a watershed moment in humankind’s history.

Acknowledgments

This work has been carried out within the framework of the National Centre of Competence in Research PlanetS supported by the Swiss National Science Foundation under grant 51NF40_205606. D.A. and S.P.Q. acknowledge the financial support of the SNSF. A.B.K. acknowledges funding provided by the University of Belgrade—Faculty of Mathematics (contract 451-03-66/2024-03/200104) through grants by the Ministry of Science, Technological Development, and Innovation of the

Republic of Serbia. E.O.G. gratefully acknowledges the financial support from the Swiss National Science Foundation (SNSF) under project grant number 200020_200399. Thank you to Geoffrey Johnson for his review and feedback during the drafting of this manuscript.

Appendix A Discussion of Mission Yields

In this appendix we give some more details on the expected yield of the future HWO and LIFE mission and explain the numbers used in the main paper. Here we also note that J. Kammerer et al. (2022) and R. Morgan et al. (2022) used slightly different models of the underlying planet occurrence rate (see Figure 5, which should be taken into account with regard to the comparison in Figure 4).

J. Kammerer et al. (2022) using their “2 m aperture at 5% throughput baseline scenario” for LIFE, derive $N_{2\text{ m/optHZ}} = 16^{+18}_{-10}$ detections of terrestrial planets in the optimistic HZ (a wider criterion than EEC) of FGK stars. Using an extended scenario with a 3.5 m aperture setup, J. Kammerer et al. (2022) this number grows to $N_{3.5\text{ m/optHZ}} = 38^{+38}_{-21}$. This number is based on the most optimistic scenario for HZ rocky planet occurrence rates in J. Kammerer et al. (2022).

For the second discussed “LIFE extended” $N_{3.5\text{ m/optHZ}} = 38$ case, the framework presented here (again, using the Jeffreys prior) would derive $\eta_{\text{obs}} = 0.012$ and a 99.9% upper limit of $\eta_{\text{obs}} < 0.13$. This means that in case of an unsuccessful search for habitable planets in this sample of 38 planets, we would be able to constrain η_{Hab} for FGK stars to be smaller than 13% with 99.9% belief.

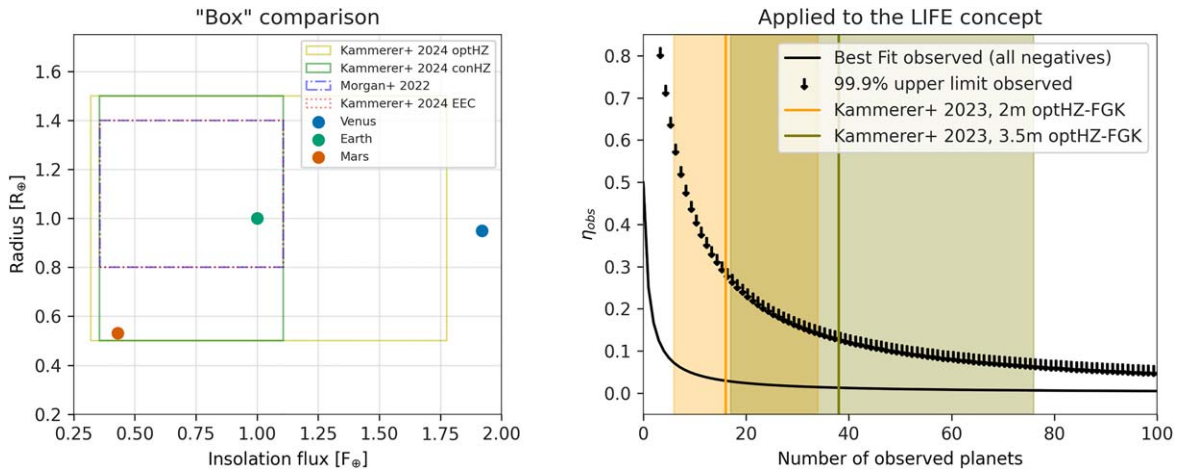


Figure 5. Left: comparison of the radius/insolation flux parameter spaces used in this paper to define the various exoplanet samples. Right: same as Figure 4 for LIFE but with a wider “optimistic habitable zone (optHZ)” sample.

Appendix B Illustrating the Priors and Difference Between Prior Choices

In this section we provide a few more figures to help the reader understand the use of the different priors in the context analyzed

here. Figure 1 compares the various priors used in this study. Figure 6 illustrated the exemplary case of the posterior distributions resulting from one positive detection in 100 (perfect) observations, assuming different priors. Figure 7 shows the difference between the results of different priors as function of the survey size.

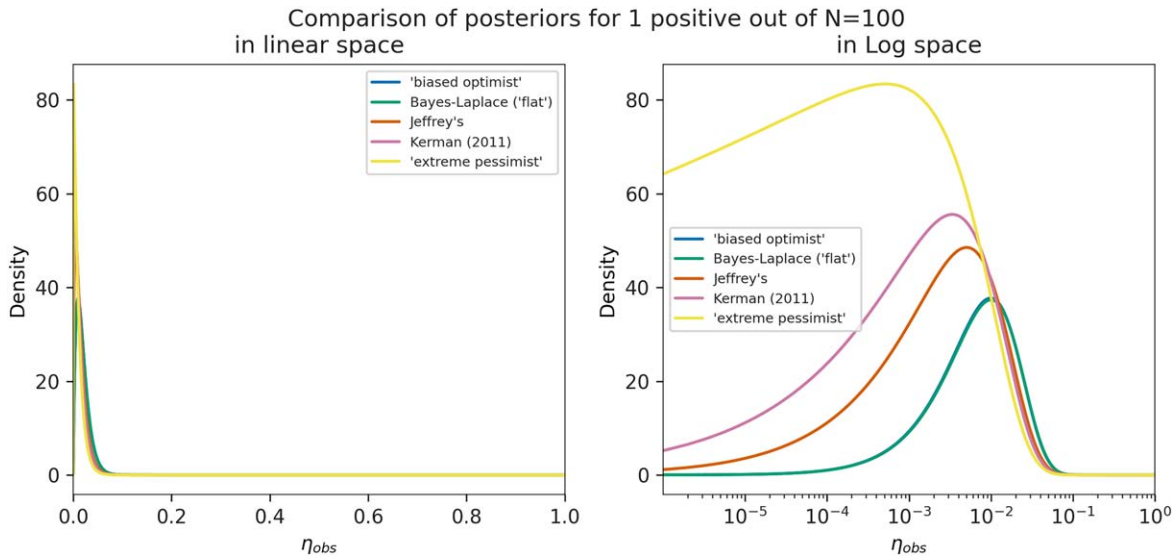


Figure 6. Example of the posterior distributions resulting from one positive detection in 100 (perfect) observations, assuming different priors. While in linear space all priors seem to constrain the resulting value around 10^{-2} (or 1%), the logarithmic plot illustrates that the more “pessimistic” priors produces more statistical weight for values of η_{obs} much smaller than 10^{-2} .

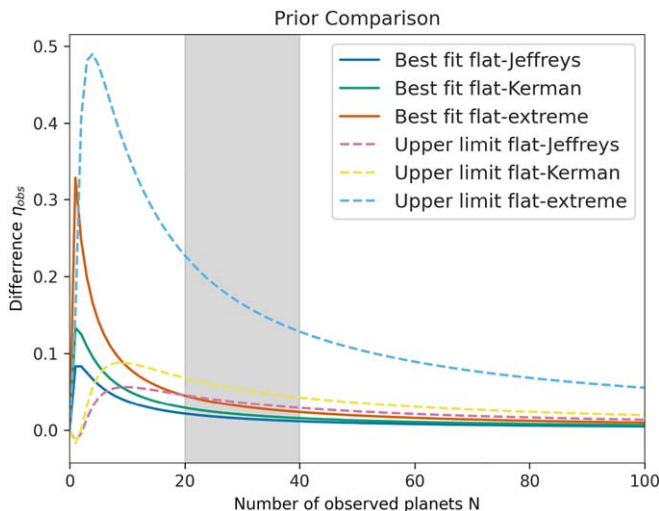


Figure 7. Comparison of results derived by this framework as function of the number of observations: differences in “best fit” and 99.9% upper limits between “optimistic” (flat) and “neutral/pessimistic” (Jeffreys/Kerman, extreme) prior analysis. $N = 20-40$ seems to be enough to be (relatively) independent of the choice of priors.

ORCID iDs

Daniel Angerhausen  <https://orcid.org/0000-0001-6138-8633>
 Amedeo Balbi  <https://orcid.org/0000-0002-3929-6932>
 Andjelka B. Kovačević  <https://orcid.org/0000-0001-5139-1978>
 Emily O. Garvin  <https://orcid.org/0000-0003-2530-9330>
 Sascha P. Quanz  <https://orcid.org/0000-0003-3829-7412>

References

- Angerhausen, D., Pidhorodetska, D., Leung, M., et al. 2024, *AJ*, 167, 128
 Balbi, A., & Grimaldi, C. 2020, *PNAS*, 117, 202007560
 Balbi, A., & Lingam, M. 2023, *MNRAS*, 522, 3117
 Benner, S. A. 2014, *OLEB*, 44, 339
 Benz, W., Broeg, C., Fortier, A., et al. 2021, *ExA*, 51, 109
 Bergsten, G. J., Pascucci, I., Mulders, G. D., Fernandes, R. B., & Koskinen, T. T. 2022, *AJ*, 164, 190
 Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Sci*, 327, 977
 Brown, L. D., Cai, T. T., & DasGupta, A. 2001, *StaSc*, 16, 101
 Bryson, S., Kunimoto, M., Kopparapu, R. K., et al. 2021, *AJ*, 161, 36
 Burn, R., Schlecker, M., Mordasini, C., et al. 2021, *A&A*, 656, A72
 Catling, D. C., Krissansen-Totton, J., Kiang, N. Y., et al. 2018, *AsBio*, 18, 709
 Chopra, A., & Lineweaver, C. H. 2016, *AsBio*, 16, 7
 Cowan, G. 2007, in ASP Conf. Ser. 371, Statistical Challenges in Modern Astronomy IV, ed. G. J. Babu & E. D. Feigelson (San Francisco, CA: ASP), 75
 Dannert, F. A., Ottiger, M., Quanz, S. P., et al. 2022, *A&A*, 664, A22
 Deamer, D. 2022, *Life*, 12, 1429
 Domagal-Goldman, S. D., Segura, A., Claire, M. W., Robinson, T. D., & Meadows, V. S. 2014, *ApJ*, 792, 90
 Drake, F. 1965, *The Radio Search for Intelligent Extraterrestrial Life* (Amsterdam: Elsevier), 323
 Emsenhuber, A., Mordasini, C., Burn, R., et al. 2021, *A&A*, 656, A70
 Fisher, T., Janin, E., & Walker, S. I. 2023, arXiv:2310.05359
 Foote, S., Sinhad, P., Mathis, C., & Walker, S. I. 2023, *AsBio*, 23, 1189
 Fujii, Y., Angerhausen, D., Deitrick, R., et al. 2018, *AsBio*, 18, 739
 Gaudi, B. S., Seager, S., Mennesson, B., et al. 2020, arXiv:2001.06683
 Gillen, C., Jeancolas, C., McMahon, S., & Vickers, P. 2023, *AsBio*, 23, 1228
 Gilmozzi, R., & Spyromilio, J. 2007, *Msngr*, 127, 11
 Green, J., Hoehler, T., Neveu, M., et al. 2021, *Natur*, 598, 575
 Greene, T. P., Line, M. R., Montero, C., et al. 2016, *ApJ*, 817, 17
 Grenfell, J. L. 2017, *PhR*, 713, 1
 Haqq-Misra, J., Kopparapu, R. K., & Schwieterman, E. 2020, *AsBio*, 20, 572
 Hardegree-Ullman, K. K., Apai, D., Haffert, S. Y., et al. 2025, *AJ*, 169, 171
 Hulson, P.-J. F., Hanselman, D. H., & Quinn, T. J. I. 2011, *ICJMS*, 69, 281
 Jin, S., Thulin, M., & Larsson, R. 2017, *Am. Stat.*, 71, 106
 Johnson, A., Ott, M., & Dogucu, M. 2022, *Bayes Rules!: An Introduction to Applied Bayesian Modeling*, Chapman & Hall/CRC Texts in Statistical Science (Boca Raton, FL: CRC Press), <https://books.google.ch/books?id=eVNcEAAAQBAJ>
 Kammerer, J., Quanz, S. P., Dannert, F. & LIFE Collaboration 2022, *A&A*, 668, A52
 Kashyap, V. L., van Dyk, D. A., Connors, A., et al. 2010, *ApJ*, 719, 900
 Kasper, M., Cerpa Urta, N., Pathak, P., et al. 2021, *Msngr*, 182, 38
 Kerman, J. 2011, *EJSta*, 5, 1450
 Kimura, T., & Ikoma, M. 2022, *NatAs*, 6, 1296
 Kipping, D. 2020, *PNAS*, 117, 11995
 Kipping, D. 2021, *PNAS*, 118, e2026808118
 Kipping, D., & Lewis, G. 2025, *IJAsB*, 24, e5
 Kite, E. S., & Ford, E. B. 2018, *ApJ*, 864, 75
 Kopparapu, R. K., Hébrard, E., Belikov, R., et al. 2018, *ApJ*, 856, 122
 Krissansen-Totton, J., Thompson, M., Galloway, M. L., & Fortney, J. J. 2022, *NatAs*, 6, 189
 Longo, A., & Damer, B. 2020, *Life*, 10, 52
 Lustig-Yaeger, J., Meadows, V. S., Tovar Mendoza, G., et al. 2018, *AJ*, 156, 301
 Meadows, V., Graham, H., Abrahamsson, V., et al. 2022, arXiv:2210.14293
 Meadows, V. S., Reinhard, C. T., Arney, G. N., et al. 2018, *AsBio*, 18, 630
 Morgan, R., Savransky, D., Turmon, M., et al. 2022, *Proc. SPIE*, 12180, 1218020
 National Academies of Sciences, Engineering, and Medicine 2021, *Pathways to Discovery in Astronomy and Astrophysics for the 2020s* (Washington, DC: The National Academies Press)
 Quanz, S. P., Absil, O., Benz, W., et al. 2022a, *ExA*, 54, 1197
 Quanz, S. P., Ottiger, M., Fontanet, E., et al. 2022b, *A&A*, 664, A21
 Ranjan, S., & Sasselov, D. D. 2016, *AsBio*, 16, 68
 Rauer, H., Catala, C., Aerts, C., et al. 2014, *ExA*, 38, 249
 Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2014, *JATIS*, 1, 014003
 Robinson, T. D., Meadows, V. S., & Crisp, D. 2010, *ApJL*, 721, L67
 Sasselov, D. D., Grotzinger, J. P., & Sutherland, J. D. 2020, *SciA*, 6, eaax3419
 Schlecker, M., Apai, D., Lichtenberg, T., et al. 2024, *PSJ*, 5, 3
 Schwieterman, E. W., Faucher, T. J., Haqq-Misra, J., et al. 2024, *ApJ*, 969, 20
 Schwieterman, E. W., Kiang, N. Y., Parenteau, M. N., et al. 2018, *AsBio*, 18, 663
 Seager, S. 2014, *PNAS of the United States of America*, 111, 12634
 Seager, S. 2018, *IJAsB*, 17, 294
 Smith, H., & Mathis, C. 2023, *BiEss*
 Spiegel, D. S., & Turner, E. L. 2012, *PNAS*, 109, 395
 Stark, C. C., Belikov, R., Bolcar, M. R., et al. 2019, *JATIS*, 5, 024009
 The LUVOIR Team 2019, arXiv:1912.06219
 Tinetti, G., Eccleston, P., Haswell, C., et al. 2021, arXiv:2104.04824
 Truitt, A. R., Young, P. A., Walker, S. I., & Spacek, A. 2020, *AJ*, 159, 55
 Vickers, P., Cowie, C., Dick, S. J., et al. 2023, *AsBio*, 23, 1202
 Walker, S. I., Bains, W., Cronin, L., et al. 2018, *AsBio*, 18, 779