The Astronomical Journal, 166:112 (9pp), 2023 September https://doi.org/10.3847/1538-3881/ace105

A Hierarchical Bayesian Framework for Inferring the Stellar Obliquity Distribution

Jiayin Dong[1] and Daniel Foreman-Mackey
Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA; jdong@flatironinstitute.org
Received 2023 May 1; revised 2023 June 6; accepted 2023 June 21; published 2023 August 17

Abstract

Stellar obliquity, the angle between a planet's orbital axis and its host star's spin axis, traces the formation and evolution of a planetary system. In transiting-exoplanet observations, only the sky-projected stellar obliquity can be measured, but this can be deprojected using an estimate of the stellar obliquity. In this paper, we introduce a flexible, hierarchical Bayesian framework that can be used to infer the stellar obliquity distribution solely from sky-projected stellar obliquities, including stellar inclination measurements when available. We demonstrate that while a constraint on the stellar inclination is crucial for measuring the obliquity of an individual system, it is not required for robust determination of the population-level stellar obliquity distribution. In practice, the constraints on the stellar obliquity distribution are mainly driven by the sky-projected stellar obliquities. When applying the framework to all systems with measured sky-projected stellar obliquity, which are mostly hot Jupiter systems, we find that the inferred population-level obliquity distribution is unimodal and peaked at zero degrees. Misaligned systems have nearly isotropic stellar obliquities with no strong clustering near 90°. The diverse range of stellar obliquities prefers dynamic mechanisms, such as planet–planet scattering after a convergent disk migration, which could produce both prograde and retrograde orbits of close-in planets with no strong inclination concentrations other than that at 0°.

Unified Astronomy Thesaurus concepts: Exoplanets (498); Bayesian statistics (1900); Exoplanet dynamics (490)

1. Introduction

Stellar obliquity ψ is the angle between a planet's orbital axis $\hat{\boldsymbol{n}}_{\rm orb}$ and the host star's spin axis $\hat{\boldsymbol{n}}_{\star}$. This angle is an important tracer of a planetary system's formation environment and dynamical evolution. The evolution of stellar obliquity can be roughly broken down into three stages. First, the formation and evolution of a protoplanetary disk determine the primordial stellar obliquity (e.g., Bate et al. 2010; Lai et al. 2011; Batygin 2012). Second, post-formation dynamical evolution in the planetary system, such as planet–planet scattering (e.g., Rasio & Ford 1996; Chatterjee et al. 2008; Nagasawa et al. 2008; Beaugé & Nesvorný 2012), von Zeipel–Lidov–Kozai (ZLK) mechanisms (e.g., Wu & Murray 2003; Naoz 2016), and secular chaos (Wu & Lithwick 2011), can excite the mutual inclinations between planetary or stellar companions and alter the stellar obliquity. Lastly, tidal force can reduce the stellar obliquity by realigning the host star's spin axis with the planet's orbital axis, if the tidal dissipation in the star is efficient (e.g., Winn et al. 2010; Albrecht et al. 2012). Additionally, massive stars with convective cores could generate internal gravity waves and dissipate angular momentum to their radiative zones, potentially affecting the stellar obliquity (Rogers et al. 2012, 2013).

It is as yet unclear if, and to what extent, all of these physical and dynamic processes apply to exoplanetary systems. These proposed mechanisms all make different predictions on stellar obliquity distributions with a focus on hot Jupiter systems (see Dawson & Johnson 2018; Albrecht et al. 2022, and references therein). For example, the secular chaos mechanism tends to produce a stellar obliquity distribution with $\psi < 90°$ (e.g., Teyssandier et al. 2019). The stellar ZLK mechanism predicts a bimodal stellar obliquity distribution, concentrated at 40° and 140° (e.g., Fabrycky & Tremaine 2007; Anderson et al. 2016; Vick et al. 2019), assuming zero stellar obliquity when the ZLK oscillation begins, or a broad stellar obliquity peaked near 90° if we drop the assumption (Vick et al. 2023). The multiple-planet scattering mechanism results in a majority of aligned systems, with a small fraction of systems at a diverse range of stellar obliquities (e.g., Beaugé & Nesvorný 2012). With these predictions in mind, we aim to determine the dominant mechanisms responsible for shaping close-in planetary systems by characterizing the stellar obliquity distribution of exoplanetary systems through a Bayesian approach.

When observing an exoplanet, typically only the sky-projected stellar obliquity λ, the angle between the projections of $\hat{\boldsymbol{n}}_{\rm orb}$ and $\hat{\boldsymbol{n}}_{\star}$ onto the plane of the sky, can be measured. This measurement is primarily obtained via the Rossiter–McLaughlin effect (McLaughlin 1924; Rossiter 1924). The stellar obliquity ψ for an individual exoplanet system can be inferred, if both the sky-projected stellar obliquity λ and the stellar inclination $i_{\star}$ are measured precisely. The relationship between ψ and $\{\lambda, i_{\star}\}$ is given by (e.g., Fabrycky & Winn 2009)

$$\cos\psi = \sin i_{\star} \sin i_{\rm orb} \cos\lambda + \cos i_{\star} \cos i_{\rm orb}, \tag{1}$$

where $i_{\rm orb}$ is the inclination angle between the vector $\boldsymbol{n}_{\rm orb}$ and the observer's line of sight, and $i_{\star}$ is the inclination angle between $\boldsymbol{n}_{\star}$ and the observer's line of sight. If an exoplanet system transits, the orbit is nearly edge-on ($i_{\rm orb} \approx 90°$), so in those cases, the relationship becomes approximately

$$\cos\psi \approx \sin i_{\star} \cos\lambda, \tag{2}$$

although this is not a simplification that we are required to make in this paper.

[1] Flatiron Research Fellow.

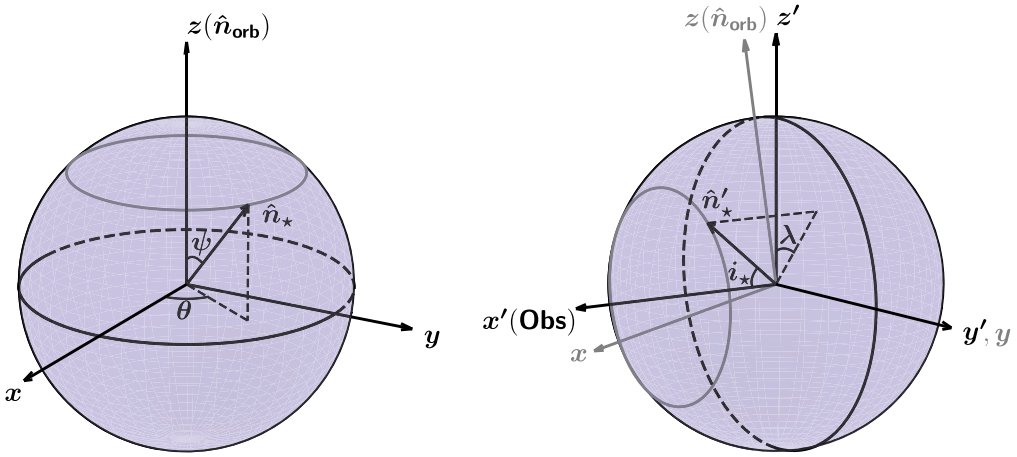

Figure 1. Two coordinate systems that describe the stellar spin axis $\hat{\boldsymbol{n}}_\star$ and the planet's orbital axis $\hat{\boldsymbol{n}}_{\rm orb}$. The $\{\psi, \theta\}$ coordinate system setup is motivated by the physical properties of a planetary system. The $\{\lambda, i_\star\}$ coordinate system setup is motivated by observables (✎).

In some cases, stellar inclinations can be constrained via, for example, photometric and spectroscopic rotational modulation introduced by starspots for cool stars (e.g., Masuda & Winn 2020; Albrecht et al. 2021), gravity darkening for fast-rotating stars (e.g., Barnes 2009; Barnes et al. 2011), and asteroseismology for bright stars (e.g., Chaplin et al. 2013). However, for the vast majority of exoplanet systems, $i_\star$ measurements are not feasible. In these cases, it is still possible to infer their stellar obliquities from the sky-projected obliquities, assuming isotropic stellar inclinations; however, the inferred ψ will have greater uncertainty than the one inferred with $i_\star$ measurement (Fabrycky & Winn 2009).

The relationship between the distributions of stellar obliquity, sky-projected stellar obliquity, and stellar inclination is still not fully understood. In this study, we aim to gain a deeper understanding of this relationship and develop a statistical approach to inferring the stellar obliquity distribution. In Section 2, we find the expression of sky-projected stellar obliquity λ and stellar inclination $i_\star$ in terms of the orbital inclination $i_{\rm orb}$, the stellar obliquity ψ, and the azimuthal angle of the stellar spin axis θ. In Section 3, we introduce a flexible, hierarchical Bayesian framework that allows us to infer the stellar obliquity distribution of a sample. In Section 4, we examine the framework with simulated data and show that the inferred stellar obliquity distribution from sky-projected stellar obliquities is robust even if the $i_\star$ information is not provided. Lastly, in Section 5, we apply the framework to real observations and derive the stellar obliquity distribution for exoplanets. We discuss the implication of the stellar obliquity distribution for hot Jupiter origins.

2. Coordinate Setup and Transformation

In this section, we find the expression of λ or $i_\star$ in terms of the stellar obliquity ψ, the azimuthal angle of the stellar spin axis relative to the orbital axis θ, and the orbital inclination $i_{\rm orb}$. In Figure 1, we introduce two coordinate systems that describe the stellar spin axis $\hat{\boldsymbol{n}}_\star$ and the planet's orbital axis $\hat{\boldsymbol{n}}_{\rm orb}$. The setup is similar to the coordinate system setup in Fabrycky & Winn (2009) but has a different definition of the azimuthal angle of the stellar spin and coordinate orientation.

The $\{\psi, \theta\}$ coordinate system, shown in the left panel in Figure 1, is designed to describe the physical properties of a system. We set the planetary orbital axis $\hat{\boldsymbol{n}}_{\rm orb}$ as the $\boldsymbol{z}$-axis. To define the stellar spin vector, we introduce the azimuthal angle of the vector around the orbital axis, θ. The stellar spin axis $\hat{\boldsymbol{n}}_\star$ can be written in terms of ψ and θ as

$$\hat{\boldsymbol{n}}_\star = \sin\psi\cos\theta\hat{\boldsymbol{x}} + \sin\psi\sin\theta\hat{\boldsymbol{y}} + \cos\psi\hat{\boldsymbol{z}}. \tag{3}$$

The $\{\lambda, i_\star\}$ coordinate system, shown in the right panel in Figure 1, is designed to describe observables of a system. We set the observer's line of sight as the $\boldsymbol{x}'$-axis. Since only the difference between the sky-projected orbital axis and the sky-projected stellar spin axis can be measured, we conveniently align the projected orbital axis with the $\boldsymbol{z}$-axis. The stellar spin axis could be written in terms of λ and $i_\star$ as

$$\hat{\boldsymbol{n}}'_\star = \cos i_\star\hat{\boldsymbol{x}}' + \sin i_\star\sin\lambda\hat{\boldsymbol{y}}' + \sin i_\star\cos\lambda\hat{\boldsymbol{z}}'. \tag{4}$$

To transform the stellar spin vector $\hat{\boldsymbol{n}}_\star$ from the $\boldsymbol{xyz}$ coordinates to the $\boldsymbol{x'y'z'}$ coordinates, we rotate the $\boldsymbol{xyz}$ coordinates by an angle $\pi/2 - i_{\rm orb}$ about the $\boldsymbol{y}'$-axis (the same as the $\boldsymbol{y}$-axis). Conventionally, the inclination angle of the orbital axis $\hat{\boldsymbol{n}}_{\rm orb}$ is assumed to be less than 90°, as we adopt here. Applying the rotation matrix such that

$$\hat{\boldsymbol{n}}'_\star = \begin{bmatrix} \sin i_{\rm orb} & 0 & \cos i_{\rm orb} \\ 0 & 1 & 0 \\ -\cos i_{\rm orb} & 0 & \sin i_{\rm orb} \end{bmatrix} \hat{\boldsymbol{n}}_\star, \tag{5}$$

we find the expression of $\hat{\boldsymbol{n}}'_\star$ in $\boldsymbol{x'y'z'}$ coordinates in terms of ψ, θ, and $i_{\rm orb}$:

$$\begin{aligned}\hat{\boldsymbol{n}}'_\star = &(\sin\psi\cos\theta\sin i_{\rm orb} + \cos\psi\cos i_{\rm orb})\hat{\boldsymbol{x}}' \\ &+ \sin\psi\sin\theta\hat{\boldsymbol{y}}' \\ &+ (-\sin\psi\cos\theta\cos i_{\rm orb} + \cos\psi\sin i_{\rm orb})\hat{\boldsymbol{z}}'.\end{aligned} \tag{6}$$

Equating Equations (6) and (4), we find the expression of λ or $i_\star$ in terms of ψ, θ, and $i_{\rm orb}$. First, from the $\hat{\boldsymbol{x}}'$ terms, we get

$$i_\star = \cos[-1](\sin\psi\cos\theta\sin i_{\rm orb} + \cos\psi\cos i_{\rm orb}). \tag{7}$$

Next, dividing the $\hat{\boldsymbol{y}}'$ terms by the $\hat{\boldsymbol{z}}'$ terms, we get

$$\lambda = \tan[-1]\left(\frac{\sin\psi\sin\theta}{-\sin\psi\cos\theta\cos i_{\rm orb} + \cos\psi\sin i_{\rm orb}}\right). \tag{8}$$

These two relations will be used in the hierarchical Bayesian framework to infer the stellar obliquity distribution.

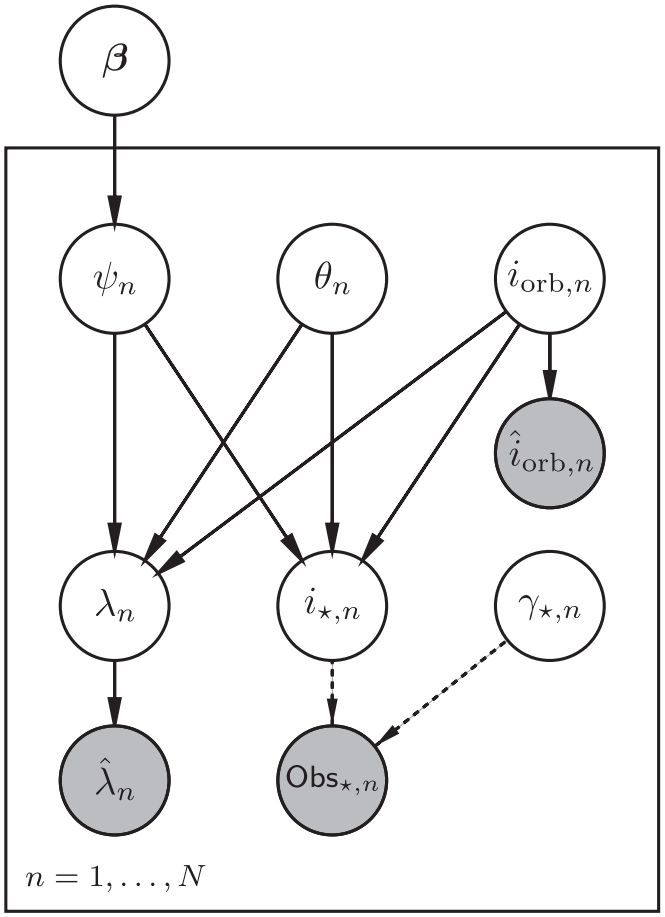

Figure 2. A probabilistic graphical model showing the conditional structure of our hierarchical Bayesian framework for inferring the stellar obliquity distribution of exoplanetary systems. The population model is parameterized by hyperparameters $\boldsymbol{\beta}$, and it is constrained by the stellar obliquity ψ_n of each n_{th} individual system. The stellar obliquity ψ_n, azimuthal angle θ_n, and orbital inclination $i_{\mathrm{orb},n}$ are constrained by the sky-projected stellar obliquity λ_n and, if available, the stellar inclination $i_{\star,n}$. Properties of the star other than its inclination, such as its radius and rotation period, are represented by $\gamma_{\star,n}$, and these can inform the constraint on $i_{\star,n}$, if available. $\hat{\lambda}_n$ comprises the measured sky-projected stellar obliquity and its uncertainty, $\mathrm{Obs}_{\star,n}$ contains all observed properties of the star with uncertainties, and $\hat{i}_{\mathrm{orb},n}$ comprises the measured orbital inclination and its uncertainty (✎).

3. Hierarchical Bayesian Framework

To find the stellar obliquity distribution of exoplanetary systems, we develop a hierarchical Bayesian framework that takes measurements of the observed sky-projected stellar obliquity λ and orbital inclination i_{orb} as input data. If a measurement of the stellar inclination $i_\star$ is available, it can be provided or inferred from the stellar rotation period P_{rot}, stellar radius $R_\star$, and sky-projected rotational broadening velocity $v \sin i_\star$. In the absence of an $i_\star$, P_{rot}, or $v \sin i_\star$ measurement, the stellar obliquity ψ distribution is inferred without $i_\star$ likelihoods.

Figure 2 illustrates the probabilistic graphical model for our hierarchical Bayesian framework. We aim to constrain a set of hyperparameters $\boldsymbol{\beta}$ that describe the stellar obliquity distribution. The parameter set $\boldsymbol{\beta}$ is constrained by N individual systems, where each ψ_n is simultaneously fit based on the entire sample of sky-projected stellar obliquities λ_n and, if available, stellar inclinations $i_{\star,n}$. The parameter $\gamma_{\star,n}$ contains all stellar properties other than $i_{\star,n}$, such as the stellar rotation period $P_{\mathrm{rot},n}$, radius $R_{\star,n}$, and projected rotational velocity $v \sin i_{\star,n}$, if they are known.

In Figure 2, the observed values of λ_n and $i_{\mathrm{orb},n}$ are indicated by $\hat{\lambda}_n$ and $\hat{i}_{\mathrm{orb},n}$, respectively. The $\hat{\lambda}_n$ measurements typically come from the Rossiter–McLaughlin effect or gravity darkening, and $\hat{i}_{\mathrm{orb},n}$ is constrained by the transit light curve.

The constraints on stellar inclination are somewhat more complicated. In Figure 2, the $\mathrm{Obs}_{\star,n}$ node includes any observed data that directly constrains $i_{\star,n}$. This could include a direct measurement of $\hat{i}_{\star,n}$ (via gravity darkening or asteroseismology; for example Barnes 2009; Chaplin et al. 2013), or a measurement of the sky-projected stellar rotational line broadening $\hat{v} \sin i_{\star,n}$, which is related to $i_\star$ as $v \sin i_\star = 2\pi R_\star / P_{\mathrm{rot}}$ (Masuda & Winn 2020).

For the stellar obliquity distribution, we model the $\cos\psi$ distribution instead of the ψ distribution to understand whether the stellar obliquity is isotropically distributed. If the stellar obliquity distribution is isotropic, $\cos\psi$ is uniformly distributed between -1 and 1. To flexibly model the stellar obliquity distribution, we select a multicomponent mixture of beta distributions with hyperparameters $\boldsymbol{\beta} = \{\boldsymbol{w}, \boldsymbol{\mu}, \boldsymbol{\kappa}\}$, where each hyperparameter has a dimension of the number of components. For example, for a two-component mixture model, each hyperparameter has a dimension of 2. We have $\boldsymbol{w} = \{w_0, w_1\}$, $\boldsymbol{\mu} = \{\mu_0, \mu_1\}$, and $\boldsymbol{\kappa} = \{\kappa_0, \kappa_1\}$. The hyperparameter $\boldsymbol{w}$ describes the weight of each component. The hyperparameters $\boldsymbol{\mu}$ and $1/\boldsymbol{\kappa}$ correspond to the mean and variance of each beta distribution component, respectively, a reparameterization suggested by Gelman et al. (2014). The greater the value of κ, the smaller the variance (i.e., the distribution is more concentrated). The relations between μ and κ and the standard α and β parameters in the beta distribution are $\alpha = \mu\kappa$ and $\beta = (1-\mu)\kappa$. This mixture distribution has the capacity to capture anything from an isotropic distribution to a strongly bimodal population. Then, under this two-component model, the probability density function for $\cos\psi$ is

$$\begin{aligned} w_{0,1} &\sim \mathrm{Bernoulli}(1/2) \\ u_{0,1} &\sim \mathrm{Beta}(\mu_{0,1}\kappa_{0,1}, (1-\mu_{0,1})\kappa_{0,1}) \\ \cos\psi &= 2(w_0 u_0 + w_1 u_1) - 1. \end{aligned} \tag{9}$$

Since the beta distribution is defined on the interval [0, 1] whereas the support of $\cos\psi$ is from -1 to 1, we extend the mixture distribution's support [0, 1] to $[-1, 1]$ using a linear transformation (i.e., $2(w_0 u_0 + w_1 u_1) - 1$). For hyperparameters μ and κ, we adopt the following priors:

$$\begin{aligned} \mu_{0,1} &\sim \mathcal{U}(0, 1) \\ \log\kappa_{0,1} &\sim \mathcal{N}(0, 3), \end{aligned} \tag{10}$$

where $\mu_{0,1}$ is uniformly distributed between 0 and 1 and $\log\kappa_{0,1}$ is normally distributed with a mean of 0 and a standard deviation of 3. To deal with label switching in the mixture model, we remove the symmetry by forcing the vector $\boldsymbol{\mu}$ to be ordered. Notably, when applying the framework to a small sample size with $N \lesssim 50$, the choice of hyperpriors for the beta distribution could impact the inferred distribution (e.g., Gelman et al. 2014; Nagpal et al. 2022). To ensure the robustness of the inferred distributions in such cases, it is crucial to test their sensitivity to different hyperpriors. We also note that our framework is flexible to customization to distributions other than the beta distribution for population-level stellar obliquity

distribution inference. For example, if instead studying the stellar obliquity distribution in the ψ–angle space, a von Mises distribution could be used to evaluate the mean value and the dispersion of the angle.

Next, the model parameter priors are the following:

$$\theta_n \sim \mathcal{U}(0, \pi)$$
$$\cos i_{\rm orb,n} \sim \mathcal{U}(0, 1). \quad (11)$$

If $\gamma_{\star,n}$ is available, we construct normal distributions with means and standard deviations from measurements. For the orbital inclination $i_{\rm orb}$, following the convention, we limit it to $[0, \pi/2]$, i.e., the orbital axis always points to us. To not underestimate ψ, we then set $i_\star$ to vary from 0° to 180° such that the stellar spin axis could either point to us or point away from us. Besides, since λ and $-\lambda$ correspond to the same ψ solution, we limit λ to $[0, \pi]$ and thus θ to $[0, \pi]$ to avoid bimodal distributions of λ and θ. We find doing so greatly improves the sampling performance while not compromising the inference of the ψ distribution due to the symmetry. The $[0, \pi]$ support avoids the otherwise-expected discontinuity at $\theta = 0$ and $\theta = \pi$.

The likelihood functions follow

$$\mathcal{L}(\lambda) \sim \prod_{n=1}^{N} \mathcal{N}(\hat{\lambda}_n, \sigma_{\hat{\lambda}_n})$$
$$\mathcal{L}(i_\star) \sim \prod_{n=1}^{N} \mathcal{N}({\rm Obs}_{\star,n}, \sigma_{{\rm Obs}_{\star,n}})({\rm optional})$$
$$\mathcal{L}(i_{\rm orb}) \sim \prod_{n=1}^{N} \mathcal{N}(\hat{i}_{\rm orb,n}, \sigma_{\hat{i}_{\rm orb,n}}). \quad (12)$$

The probabilistic model is constructed using the `PyMC` package `v5.1.2` (Wiecki et al. 2022), and the posteriors are sampled with the No-U-Turn Sampler (Hoffman & Gelman 2011), which is a gradient-based Markov Chain Monte Carlo sampling algorithm. This paper's figures and simulations are completely reproducible and are created using the `showyourwork` package. The open-source code is available on GitHub.[2]

4. Model Performance

4.1. Simulated Data

To investigate the performance of our hierarchical Bayesian framework, we apply it to simulated data, generated with a known ground-truth stellar obliquity distribution. We test the following five $\cos\psi$ distributions: a uniform distribution bounded between -1 and 1 such that $\cos\psi \sim \mathcal{U}(-1, 1)$, a beta distribution such that $\cos\psi \sim {\rm Beta}(3, 6)$ with the support being extended linearly from $[0, 1]$ to $[-1, 1]$, and three truncated normal distributions bounded between -1 and 1, such that $\cos\psi \sim \mathcal{N}(0, 0.2)$, $\cos\psi \sim \mathcal{N}(-0.4, 0.2)$, and $\cos\psi \sim \mathcal{N}(0.4, 0.2)$, where the first number is the mean of the distribution and the second number is the standard deviation. For each $\cos\psi$ distribution, we randomly generate 400 samples of sky-projected stellar obliquity λ and stellar inclination $i_\star$. We assume the stellar spin axis is uniformly distributed around the planetary orbital axis in the azimuthal direction and the orbital inclination is 90°. The sampled $i_\star$ and λ here are *true* values. To simulate the observation process, we add Gaussian noise to the *true* λ and $i_\star$, using uncertainties of $\sigma_\lambda = 8°$ and $\sigma_{i_\star} = 10°$, which are typical of the literature sample (Albrecht et al. 2022). Using these simulated λ measurements and their uncertainties, we infer the $\cos\psi$ distribution of the sample with both λ and $i_\star$ likelihoods or λ likelihood only.

In Figure 3, we present the results of this experiment, plotting the inferred stellar obliquity distributions, compared to the ground-truth distributions. Since the simulated stellar obliquity distributions only have a single component, we model the data with a single beta distribution. Each row of Figure 3 corresponds to a different simulation distribution. The blue curve and contours in the left column are the median and 1σ and 2σ uncertainties of the inferred $\cos\psi$ distribution when constraints on stellar inclination are not included. The right column shows the same inferences (in orange) with stellar inclination information. Surprisingly, Figure 3 demonstrates that our inference procedure recovers the true distribution for $\cos\psi$ equally well, regardless of the inclusion of $i_\star$ measurements. Despite the fact that the inferred distributions without $i_\star$ measurements have marginally wider uncertainties, as indicated by the shallow color contours, the modes and widths of the inferred stellar obliquity distributions are consistent with or without $i_\star$ likelihood. Since the injected distributions in rows 2–4 are normal distributions, it should not be surprising that the inferred distributions, which are beta distributions, may not exactly match the injected distributions. Given that the true distributions are not included within the support of the distributions we use to fit, it is impressive how well the underlying distributions could be recovered.

We also examine the role of orbital inclination $i_{\rm orb}$ in the stellar obliquity distribution inference. Since our study focuses on transiting-exoplanet systems, we consider an isotropic orbital inclination distribution between 80° and 90°. This broad range of inclinations corresponds to an impact parameter range from 0 to 1 with a planet–star separation $a/R_\star$ of 6. We compare the stellar obliquity distributions obtained by approximating $i_{\rm orb}$ to 90° with the distributions obtained using the actual $i_{\rm orb}$. We find the difference between the two distributions is negligible. This suggests that for transiting-exoplanet systems, approximating orbital inclinations as 90° will not compromise the stellar obliquity distribution inference.

We demonstrate through simulations that the inferred stellar obliquity distribution is robust even if the $i_\star$ information is not provided, and the $i_\star$ measurement only improves the constraint on the stellar obliquity distribution.

4.2. Jacobian Transformations between the ψ, λ, and $i_\star$ Distributions

Next, we aim to gain insight into the reasons behind the predominant role of the sky-projected stellar obliquity distribution and the less significant impact of the stellar inclination distribution in the inference of the stellar obliquity distribution. To simplify the problem, the orbital inclination of the transiting planet is assumed to be 90° in this illustration. As discussed earlier, the assumption will not compromise the stellar obliquity distribution inference for transiting planets. We pair the Cartesian components in Equations (4) and (6),

[2] https://github.com/jiayindong/obliquity//tree/24e5afe84650a8b2c5a39e2c9f71ab52edbf0043/

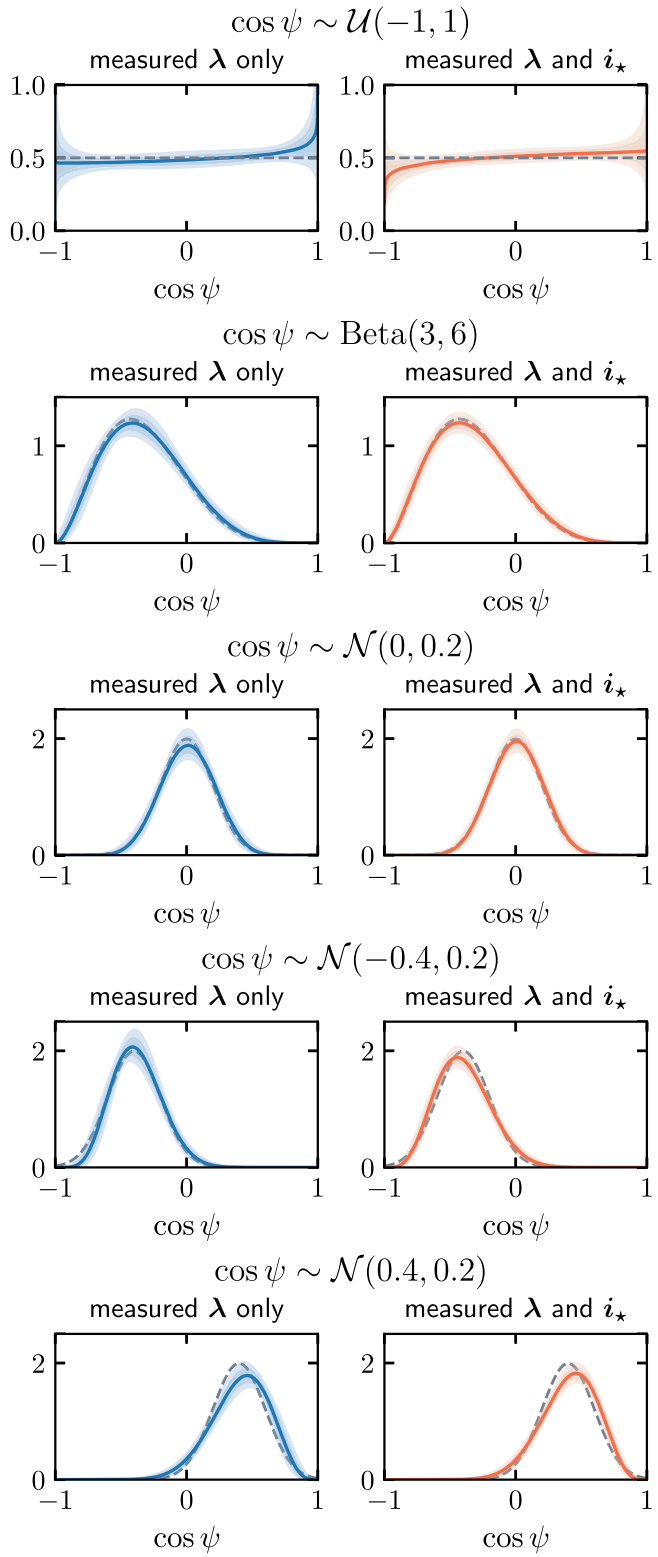

Figure 3. Inferred stellar obliquity distributions from sky-projected stellar obliquities without and with information on the stellar inclination, depicted in blue and orange curves, respectively. Each row presents a set of simulated data with the true distribution of $\cos\psi$ indicated by gray dashed curves. The shallow contours represent the 1σ and 2σ uncertainties of the inferred distribution (✎).

assuming $i_{\rm orb} = 90°$, and get

$$\sin\psi\cos\theta = \cos i_\star \tag{13}$$

$$\sin\psi\sin\theta = \sin\lambda\sin i_\star \tag{14}$$

$$\cos\psi = \cos\lambda\sin i_\star. \tag{15}$$

First, we derive the λ distribution for a given $\cos\psi$ distribution. We could find the distribution of $\cos\lambda$ using the Jacobian transformation from $\cos\psi$ and $\cos\theta$. Since ψ and θ are assumed to be independent variables, we could marginalize over θ to find the relation between the probability density functions between λ and ψ. The Jacobian transformation follows

$$p(\cos\lambda) = \int \left|\frac{\partial\cos\psi}{\partial\cos\lambda}\right| p(\cos\psi)p(\cos\theta)d\cos\theta. \tag{16}$$

Here $p(x)$ denotes the probability density distribution of x. To find ψ in terms of λ and θ, we replace $\sin i_\star$ in Equation (15) using Equation (13) and find

$$\cos\lambda = \frac{\cos\psi}{\sqrt{1 - (1 - \cos^2\psi)\cos^2\theta}}. \tag{17}$$

Reorganizing the equation, we get

$$\cos^2\psi = \frac{\cos^2\lambda\cos^2\theta - \cos^2\lambda}{\cos^2\lambda\cos^2\theta - 1} \tag{18}$$

and the partial derivative

$$\left|\frac{\partial\cos\psi}{\partial\cos\lambda}\right| = \frac{(1 - \cos^2\theta)^{1/2}}{(1 - \cos^2\theta\cos^2\lambda)^{3/2}}. \tag{19}$$

Since θ is uniformly distributed between 0 and π, $p(\cos\theta) = p(\theta)|d\theta d\cos\theta| = 1/\pi/(1 - \cos^2\theta)^{1/2}$. Putting all the parts together, we get

$$p\left(\cos\lambda\right) = \frac{1}{\pi}\int_{-1}^{1}(1 - \cos^2\theta\cos^2\lambda\Big)^{-3/2} p(\cos\psi)d\cos\theta. \tag{20}$$

In the special case where $\cos\psi$ is uniformly distributed, i.e., $p(\cos\psi) = 1/2$, Equation (20) becomes $p(\cos\lambda) = 1/\pi/\sqrt{1 - \cos^2\lambda}$, which is equivalent to $\lambda \sim \mathcal{U}(0, \pi)$. This suggests λ is uniformly distributed for an isotropic ψ distribution, as expected.

Next, we derive the $i_\star$ distribution for a given $\cos\psi$ distribution. Similarly, we first find the Jacobian transformation of $i_\star$ from ψ and θ and then marginalize over θ. It is easier to work on $\cos i_\star$ than on $i_\star$:

$$p(\cos i_\star) = \int \left|\frac{\partial\sin\psi}{\partial\cos i_\star}\right| p(\sin\psi)p(\cos\theta)d\cos\theta. \tag{21}$$

From Equation (13), we get

$$\sin\psi = \frac{\cos i_\star}{\cos\theta} \tag{22}$$

and the partial derivative

$$\left|\frac{\partial\sin\psi}{\partial\cos i_\star}\right| = \frac{1}{\cos\theta}. \tag{23}$$

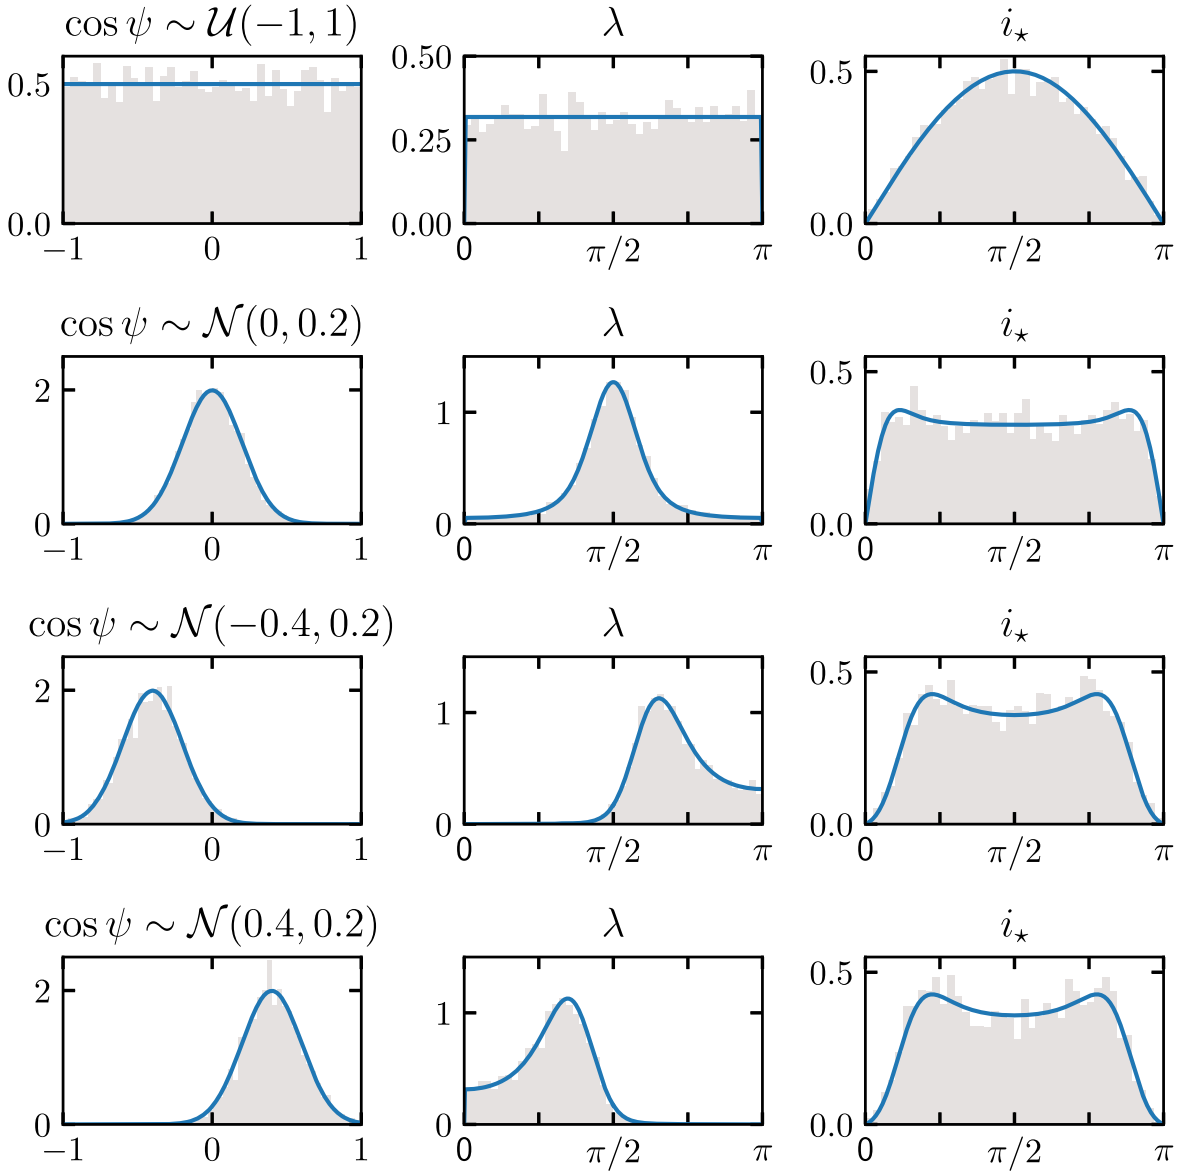

Figure 4. Simulated $\cos\psi$ distributions (first column) and the corresponding distributions of sky-projected stellar obliquity λ (second column) and stellar inclination $i_\star$ (third column). The gray histograms present the random samplings of λ and $i_\star$ from the $\cos\psi$ distributions, and the blue curves present the numerical solutions (✎).

Again, we assume θ is uniformly distributed, and this gives $p(\cos\theta) = 1/\pi/(1 - \cos^2\theta)^{1/2}$. Lastly, we transform $p(\sin\psi)$ to $p(\cos\psi)$:

$$p(\sin\psi) = \frac{2p(\cos\psi)\sin\psi}{\sqrt{1 - \sin^2\psi}}. \tag{24}$$

The factor of 2 is from two solutions of $\cos\psi$ to $\cos^2\psi = 1 - \sin^2\psi$. Combining all the pieces together, we get

$$p(\cos i_\star) = \frac{2}{\pi}\int_{\cos i_\star}^{1,-1} \frac{\cos i_\star/\cos\theta}{\sqrt{\cos^2\theta - \cos^2 i_\star}} \frac{1}{\sqrt{1 - \cos^2\theta}} p(\cos\psi) d\cos\theta, \tag{25}$$

where the integral is from $\cos i_\star$ to 1 for $\cos\theta > 0$, and from $\cos i_\star$ to -1 for $\cos\theta \leqslant 0$. Note that the lower limit of the integral is $\cos i_\star$ instead of 0 since $|\cos\theta/\cos i_\star| \geqslant 1$. If $\cos\psi$ is uniformly distributed, i.e., $p(\cos\psi) = 1/2$, the integral gives 1, which suggests $\cos i_\star$ is uniformly distributed, as expected.

Using Equations (20) and (25), we can now derive the λ and $i_\star$ distributions for any given ψ distributions, assuming the azimuthal angle of the stellar spin axis θ is random. In Figure 4, we present numerical solutions of the λ and $i_\star$ distributions for the four different $\cos\psi$ distributions used in Section 4.1. The top row of Figure 4 shows an isotropic ψ distribution, where $\cos\psi \sim \mathcal{U}(-1, 1)$. The second, third, and fourth rows of Figure 4 present truncated normal distributions of $\cos\psi$ following $\mathcal{N}(0, 0.2)$, $\mathcal{N}(-0.4, 0.2)$, and $\mathcal{N}(0.4, 0.2)$, respectively. The blue curves in each row show the numerical solutions of the λ and $i_\star$ distributions, while the gray histograms show the sampling of λ and $i_\star$ from the $\cos\psi$ and θ distributions. For a uniform $\cos\psi$ distribution, the λ distribution is uniform, and the $i_\star$ distribution is isotropic, proportional to $\sin i_\star$, as expected.

Interestingly, the λ distribution is closely related to and sensitive to the underlying ψ distribution, as demonstrated in the first and second columns in Figure 4. For different stellar obliquities, the λ distributions are distinguishable, making it

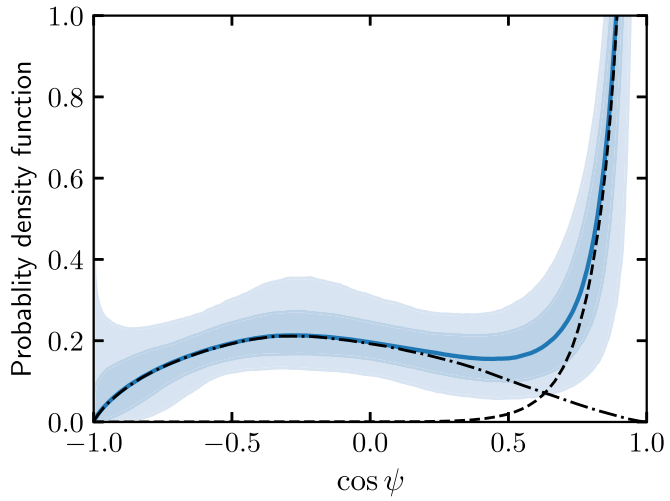

Figure 5. Inferred stellar obliquity distribution for all exoplanetary systems with sky-projected stellar obliquity measurements. This inference is based purely on the observed sky-projected obliquities using a two-component model. The blue line and shallow contours represent the median and 1σ and 2σ uncertainties of the inferred distributions. The medians of the two components are shown as dashed and dotted–dashed lines (✎).

possible to infer the ψ distribution from the λ distribution. On the other hand, the $i_\star$ distributions are less dependent on the underlying ψ distribution. Compared to an isotropic $i_\star$ distribution, the curvature of the $i_\star$ distributions for different ψ distributions differs the most at low $i_\star$ values (i.e., $i_\star < \pi/4$), which places a challenge to observational detections. Additionally, the degeneracy of the solution could be a significant issue when attempting to infer the ψ distribution from the $i_\star$ distribution. For example, when $\cos\psi \sim \mathcal{N}(-0.4, 0.2)$ or $\cos\psi \sim \mathcal{N}(0.4, 0.2)$, two $i_\star$ distributions are exactly the same, since the corresponding $i_\star$ distributions are identical for ψ distributions symmetric around $\psi = \pi/2$.

The ψ distribution can be inferred from the λ distribution without loss of information due to the strong dependency of the λ distribution on the ψ distribution. It is also worthwhile to note that although we could find a mathematical expression of ψ with λ and θ, the ψ distribution cannot be inferred from the two variables since they are not independent variables.

5. Application to Exoplanetary Systems

We next apply our hierarchical Bayesian framework to a sample of 161 exoplanetary systems with sky-projected stellar obliquity measurements, primarily consisting of hot Jupiter systems, as summarized in Albrecht et al. (2022, Table A1). We use a two-component model here for its flexibility in describing both an aligned-system population and a misaligned-system population that may or may not be concentrated at a certain angle. The inferred $\cos\psi$ distribution is shown in Figure 5. The $\cos\psi$ distribution is peaked at 1, with nearly flat behavior between -0.75 and 0.75 and no significant clustering. The distribution suggests that there is a pileup of planetary systems with stellar obliquities less than 40° and an isotropic distribution for obliquities between 40° and 140°. The fraction of aligned systems w_1 dominates the distribution, with $w_1 = 0.719 \pm 0.085$. The corresponding posteriors for the parameters of the population's beta distribution are $\mu_1 = 0.976 \pm 0.022$ and $\kappa_1 = 14.1 \pm 26.6$. On the other hand, the fraction of misaligned systems is estimated to have $w_0 = 0.281 \pm 0.085$, with posteriors of $\mu_1 = 0.434 \pm 0.088$ and $\kappa_1 = 4.2 \pm 5.6$. Note that $\mu_{0,1}$ needs to be transformed by $2\mu - 1$ to represent the true means of the $\cos\psi$ distributions.

The discrepancy between this inference for the full sample of exoplanet systems and the previous analysis of the subsample with $i_\star$ measurements (Albrecht et al. 2021) warrants further investigation. The earlier study identified a concentration of perpendicular planets and disfavored an isotropic stellar obliquity distribution. Indeed, when we apply our framework to this subsample, a concentration near $\cos\psi = -0.2$ is found in models including or excluding the stellar inclination information, as shown in Figure 6. There are at least two potential explanations for this difference: (1) the subsample with $i_\star$ measurements is small and only includes about 20 misaligned systems, and (2) the requirement for $i_\star$ measurements could introduce selection biases in the sample. First, since the sample size of the misaligned sample with $i_\star$ measurements is relatively small ($N < 30$), the observed sample, even if it is unbiased, may not be able to represent the underlying distribution. The small sample size leads to a large uncertainty on the inferred stellar obliquity distribution, shown by the 1σ and 2σ contours in Figure 6. Second, the requirement of $i_\star$ measurements could introduce selection biases. The rotation modulation technique is most applicable to cool stars with spots, but misaligned hot Jupiters are mostly found around hot stars. To measure the stellar inclinations of the hot-star hosts of misaligned hot Jupiters, the gravity darkening technique is commonly used, but the technique is biased toward detecting polar-orbit planets.

The inferred stellar obliquity distribution indicates that approximately $72\% \pm 9\%$ of the systems have a stellar obliquity of less than 40°, and approximately $28\% \pm 9\%$ of the systems follow a nearly isotropic stellar obliquity distribution ranging from ∼40° to ∼140°. These findings could have significant implications for the formation and evolution of close-in planetary systems, particularly of hot Jupiters. The diverse distribution disfavors dynamical mechanisms, such as secular chaos, which tends to produce stellar obliquities less than 90°, or stellar Kozai, which tends to produce stellar obliquities clustered at certain angles. The broad distribution of misaligned systems is in good agreement with the predicted outcome of multiple giant planets scattering after a convergent disk migration, as proposed by various studies, such as Nagasawa & Ida (2011) and Beaugé & Nesvorný (2012). The intriguing result should be further examined with a more carefully selected sample of hot Jupiters and provides opportunities to place constraints on the origin channels of hot Jupiters.

6. Summary and Discussion

In this work, we demonstrate that the stellar obliquity distribution could be robustly inferred from sky-projected stellar obliquities purely. We introduce a flexible, hierarchical Bayesian framework for stellar obliquity distribution inference. Stellar inclination measurements are optional input in the model, and if not available, they are assumed to be isotropically distributed. Our open-source hierarchical Bayesian model, available on GitHub[3], can be customized to different stellar obliquity distributions and priors for specific target samples.

[3] https://github.com/jiayindong/obliquity

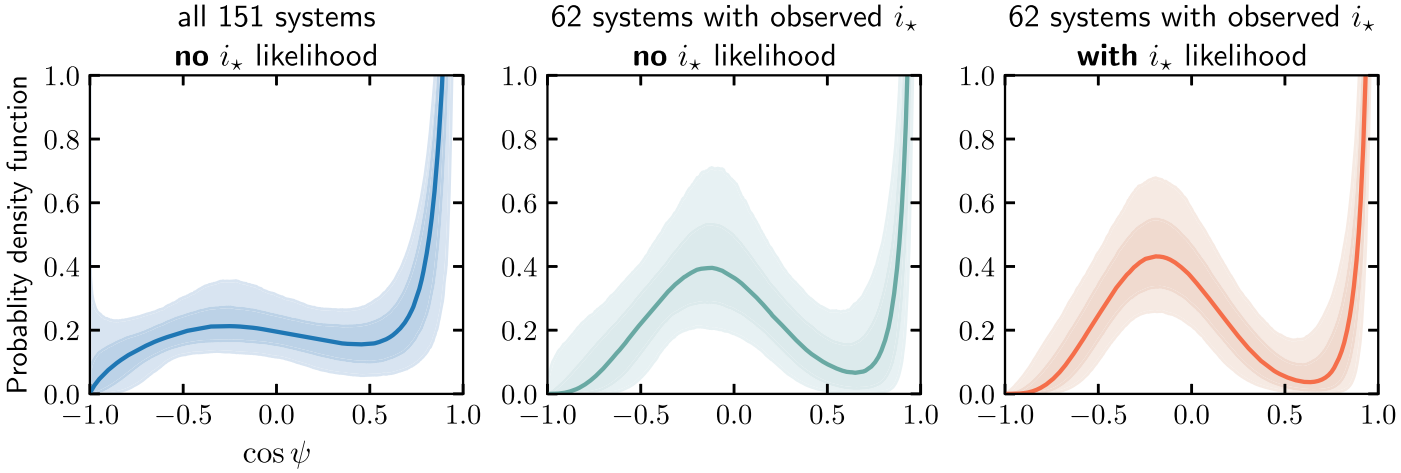

Figure 6. Inferred stellar obliquity distribution for the full sample of exoplanetary systems with sky-projected stellar obliquity measurements (left panel) and the subsample of systems with both sky-projected stellar obliquity measurements and stellar inclination measurements (middle and right panels). The subsample inference made purely from measured projected stellar obliquity is shown in the middle panel, and the subsample inference made from both measured projected stellar obliquity and stellar inclination is shown in the right panel. The line and shallow contours represent the median and 1σ and 2σ uncertainties of the inferred distributions. The samples are extracted from Albrecht et al. (2021, 2022) (✎).

It is crucial to consider the representativeness of the $i_\star$ sample when jointly modeling the stellar obliquity distribution from two data sets, one with and one without $i_\star$ measurements. An unrepresentative $i_\star$ sample could tighten the constraints on the stellar obliquities and bias the interpretation of the overall distribution.

Finally, we apply the framework to all exoplanetary systems with available sky-projected stellar obliquities and find that approximately $72\% \pm 9\%$ of the systems have a stellar obliquity less than 40°, and approximately $28\% \pm 9\%$ of the systems follow a nearly isotropic stellar obliquity distribution between ∼40° and ∼140°. The distribution could have important implications for the formation and evolution of close-in planetary systems and is worth further investigation.

Acknowledgments

We are grateful to Josh Winn and Jared Siegel for their valuable insights and feedback on this project.

This study was conducted using the showyourwork[4] reproducible workflow (Luger et al. 2021), which leverages continuous integration to automate data retrieval[5], figure generation, and manuscript compilation. The script used to produce each figure can be accessed via a link in the corresponding figure caption, as it corresponds to the latest build of the manuscript. The Git repository for this study, which includes the Jupyter notebook demonstration and case studies of the stellar obliquity distribution inference, is publicly accessible[6].

Software: `ArviZ` (Kumar et al. 2019), `Jupyter` (Kluyver et al. 2016), `Matplotlib` (Hunter 2007; Droettboom et al. 2016), `NumPy` (van der Walt et al. 2011; Harris et al. 2020), `pandas` (McKinney 2010; pandas development team 2020), `PyMC` (Wiecki et al. 2022), `SciPy` (Virtanen et al. 2020), `showyourwork` (Luger et al. 2021).

[4] https://github.com/showyourwork/showyourwork
[5] https://zenodo.org/
[6] https://github.com/jiayindong/obliquity

ORCID iDs

Jiayin Dong https://orcid.org/0000-0002-3610-6953
Daniel Foreman-Mackey https://orcid.org/0000-0002-9328-5652

References

Albrecht, S., Winn, J. N., Johnson, J. A., et al. 2012, ApJ, 757, 18
Albrecht, S. H., Dawson, R. I., & Winn, J. N. 2022, PASP, 134, 082001
Albrecht, S. H., Marcussen, M. L., Winn, J. N., Dawson, R. I., & Knudstrup, E. 2021, ApJL, 916, L1
Anderson, K. R., Storch, N. I., & Lai, D. 2016, MNRAS, 456, 3671
Barnes, J. W. 2009, ApJ, 705, 683
Barnes, J. W., Linscott, E., & Shporer, A. 2011, ApJS, 197, 10
Bate, M. R., Lodato, G., & Pringle, J. E. 2010, MNRAS, 401, 1505
Batygin, K. 2012, Natur, 491, 418
Beaugé, C., & Nesvorný, D. 2012, ApJ, 751, 119
Chaplin, W. J., Sanchis-Ojeda, R., Campante, T. L., et al. 2013, ApJ, 766, 101
Chatterjee, S., Ford, E. B., Matsumura, S., & Rasio, F. A. 2008, ApJ, 686, 580
Dawson, R. I., & Johnson, J. A. 2018, ARA&A, 56, 175
Droettboom, M., Hunter, J., Caswell, T. A., et al. 2016, Matplotlib: Matplotlib, v1.5.1, Zenodo, doi:10.5281/zenodo.44579
Fabrycky, D., & Tremaine, S. 2007, ApJ, 669, 1298
Fabrycky, D. C., & Winn, J. N. 2009, ApJ, 696, 1230
Gelman, A., Carlin, J. B., Stern, H. S., et al. 2014, Bayesian Data Analysis (London: Chapman & Hall)
Harris, C. R., Jarrod Millman, K., van der Walt, S. J., et al. 2020, Natur, 585, 357
Hoffman, M. D., & Gelman, A. 2011, arXiv:1111.4246
Hunter, J. D. 2007, CSE, 9, 90
Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in Positioning and Power in Academic Publishing: Players, Agents and Agendas, ed. F. Loizides & B. Scmidt (Amsterdam: IOS Press), 87, https://eprints.soton.ac.uk/403913/
Kumar, R., Carroll, C., Hartikainen, A., & Martin, O. 2019, JOSS, 4, 1143
Lai, D., Foucart, F., & Lin, D. N. C. 2011, MNRAS, 412, 2790
Luger, R., Bedell, M., Foreman-Mackey, D., et al. 2021, arXiv:2110.06271
Masuda, K., & Winn, J. N. 2020, AJ, 159, 81
McKinney, W. 2010, in Proc. 9th Python Science Conf., ed. S. van der Walt & J. Millman (Austin, TX: SciPy), 56
McLaughlin, D. B. 1924, ApJ, 60, 22
Nagasawa, M., & Ida, S. 2011, ApJ, 742, 72
Nagasawa, M., Ida, S., & Bessho, T. 2008, ApJ, 678, 498
Nagpal, V., Blunt, S., Bowler, B. P., et al. 2022, arXiv:2211.02121
Naoz, S. 2016, ARA&A, 54, 441
pandas development team, T 2020, pandas-dev/pandas: Pandas, latest, Zenodo, doi:10.5281/zenodo.3509134

Rasio, F. A., & Ford, E. B. 1996, Sci, 274, 954
Rogers, T. M., Lin, D. N. C., & Lau, H. H. B. 2012, ApJL, 758, L6
Rogers, T. M., Lin, D. N. C., McElwaine, J. N., & Lau, H. H. B. 2013, ApJ, 772, 21
Rossiter, R. A. 1924, ApJ, 60, 15
Teyssandier, J., Lai, D., & Vick, M. 2019, MNRAS, 486, 2265
van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22
Vick, M., Lai, D., & Anderson, K. R. 2019, MNRAS, 484, 5645
Vick, M., Su, Y., & Lai, D. 2023, ApJL, 943, L13
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261
Wiecki, T., Salvatier, J., Patil, A., et al. 2022, pymc-devs/pymc:, v4.1.7, Zenodo, doi:10.5281/zenodo.7467113
Winn, J. N., Fabrycky, D., Albrecht, S., & Johnson, J. A. 2010, ApJL, 718, L145
Wu, Y., & Lithwick, Y. 2011, ApJ, 735, 109
Wu, Y., & Murray, N. 2003, ApJ, 589, 605