The Astrophysical Journal, 808:173 (16pp), 2015 August 1 doi:10.1088/0004-637X/808/2/173

SIMULATED PHOTOEVAPORATIVE MASS LOSS FROM HOT JUPITERS IN 3D

Anjali Tripathi[1], Kaitlin M. Kratter[2], Ruth A. Murray-Clay[1,3], and Mark R. Krumholz[4]
[1] Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA; atripathi@cfa.harvard.edu
[2] Department of Astronomy and Steward Observatory, University of Arizona, Tucson, AZ 85721, USA
[3] Department of Physics, University of California, Santa Barbara, CA 93106, USA
[4] Department of Astronomy and Astrophysics, University of California, 1156 High Street, Santa Cruz, CA 95064, USA
Received 2015 May 4; accepted 2015 June 18; published 2015 July 30

ABSTRACT

Ionizing stellar photons heat the upper regions of planetary atmospheres, driving atmospheric mass loss. Gas escaping from several hot, hydrogen-rich planets has been detected using UV and X-ray transmission spectroscopy. Because these planets are tidally locked, and thus asymmetrically irradiated, escaping gas is unlikely to be spherically symmetric. In this paper, we focus on the effects of asymmetric heating on local outflow structure. We use the Athena code for hydrodynamics to produce 3D simulations of hot Jupiter mass loss that jointly model wind launching and stellar heating via photoionization. Our fiducial planet is an inflated, hot Jupiter with radius $R_p = 2.14R_{Jup}$ and mass $M_p = 0.53M_{Jup}$. We irradiate the initially neutral, atomic hydrogen atmosphere with 13.6 eV photons and compute the outflow's ionization structure. There are clear asymmetries in the atmospheric outflow, including a neutral shadow on the planet's nightside. Given an incident ionizing UV flux comparable to that of the Sun, we find a steady-state mass loss rate of $\sim 2 \times 10^{10}$ g s^{-1}. The total mass loss rate and the outflow substructure along the substellar ray show good agreement with earlier 1D models, for two different fluxes. Our 3D data cube can be used to generate the outflow's extinction spectrum during transit. As a proof of concept, we find absorption of stellar Lyα at Doppler-shifted velocities of up to ± 50 km s^{-1}. Our work provides a starting point for further 3D models that can be used to predict observable signatures of hot Jupiter mass loss.

Key words: hydrodynamics – planet–star interactions – planets and satellites: atmospheres – planets and satellites: gaseous planets

1. INTRODUCTION

Photoionization by high energy stellar radiation heats the upper layers of planetary atmospheres. This heating drives thermal mass loss and can thus play a substantial role in a planet's atmospheric evolution. Strongly irradiated hydrogen-rich planets are most susceptible to mass loss, and transit observations of close-in giant planets at UV and X-ray wavelengths have revealed atmospheric escape. However, models are required to translate these observations into constraints on mass loss rates or outflow structure. While models have been developed for hot Jupiter mass loss, none have yet consistently modeled the heating and three-dimensional (3D) structure of atmospheric escape. The inherent asymmetry in the physics of atmospheric escape, especially asymmetric irradiation expected of tidally locked hot Jupiters, necessitates full 3D modeling. To examine how asymmetric heating affects the structure of the outflow near the planet, we develop a 3D, self-consistent model of mass loss driven by photoionization heating.

The first indication of hot Jupiter mass loss came from the transmission spectroscopy of Vidal-Madjar et al. (2003), who observed decreased stellar Lyα emission during transits of the hot Jupiter, HD 209458b. The depth of the transit in Lyα, 15% ± 4%, was ten times larger than the optical transit depth of 1.5% (Charbonneau et al. 2000; Henry et al. 2000), suggesting absorption by a neutral hydrogen atmosphere larger than the planet's Roche lobe. Absorption out to Doppler equivalent velocities of ± 100 km s^{-1} indicated that the neutral gas either moved at high velocities or had a large column depth, enhancing the line's naturally broadened wings. Subsequent studies have confirmed stellar Lyα absorption by extended gas around HD 209458b (Ehrenreich et al. 2008) and the hot Jupiter HD 189733b (Lecavelier des Etangs et al. 2010), with temporal variations in the measured absorption (Ben-Jaffel 2008; Lecavelier des Etangs et al. 2012). Absorption in C II, O I, Si III, and Mg II for HD 209458b (Vidal-Madjar et al. 2004; Linsky et al. 2010; but see Ballester & Ben-Jaffel 2015) and in X-rays for HD 189733b (Poppenhaeger et al. 2013) suggest that the atmospheric outflows of these planets are metal enriched, although interpretation of the X-ray signal may be complicated by stellar variability (Llama & Shkolnik 2015). Metal line absorption has also been detected for the hot Jupiter WASP-12b (Haswell et al. 2012), which is thought to be overflowing its Roche lobe (Li et al. 2010). Stellar Lyα absorption observations of a transiting hot Neptune, GJ 436b, suggest it has an asymmetric outflow structure (Kulow et al. 2014; Ehrenreich et al. 2015).

Stellar heating induces planetary mass loss through hydrodynamic escape. This heating is primarily due to extreme UV stellar radiation, which photoionizes neutrals in the planet's upper atmosphere, liberating electrons that heat the gas through collisions. This photoionization heating creates pressure gradients that accelerate the gas from subsonic to supersonic velocities. Consequently, heated gas moves outward in the form of a planetary wind, similar to the structure of the solar wind (Parker 1958). Gas exceeding the escape velocity or displaced beyond the Roche lobe becomes unbound and escapes into space. For the large ionizing fluxes received by hot Jupiters, hydrodynamic escape is the most efficient mass loss process. Other mass loss mechanisms, including Jeans escape (Chamberlain 1963), are less efficient because they operate on individual particles, rather than a collective fluid.

One-dimensional (1D) hydrodynamic models of energy-limited escape from hot Jupiters were first calculated by Lammer et al. (2003) and Baraffe et al. (2004) to explain the Vidal-Madjar et al. (2003) observations of HD 209458b. Based on work done for the early Earth and Venus (Gross 1972; Watson et al. 1981), these pioneering calculations assumed that all energy deposited by stellar radiation goes into heating. The heat is conducted to lower radii and drives gas to larger radii, where it can absorb more flux, thus enhancing the outflow. These models predicted catastrophic evaporation for hot Jupiters. Other energy-limited models (including Hubbard et al. 2007) and more detailed 1D models that have accounted for chemistry, heating and cooling, tidal gravity, and the stellar wind (Yelle 2004; García Muñoz 2007; Murray-Clay et al. 2009, hereafter MC09) have suggested less dramatic mass loss for hot Jupiters. Hot Jupiters, unlike the early Earth and Venus, do not efficiently conduct heat downward (García Muñoz 2007; MC09).

Not modeled by 1D studies are the inherent asymmetries in atmospheric escape processes. The stellar wind's pressure confinement, rotation from the Coriolis force, and magnetic fields add to the asymmetry of escaping atmospheric gas. Day and night differences due to the stellar wind have been captured by mass loss models in 2D (Stone & Proga 2009; Tremblin & Chiang 2013) and in 3D (Bisikalo et al. 2013; Llama et al. 2013). These and other 3D simulations (Schneiter et al. 2007; Cohen et al. 2011; Bourrier & Lecavelier des Etangs 2013; Matsakos et al. 2015), which variously include orbital motion or magnetic fields, have not directly included the planetary wind's production by ionizing radiation from the host star. Instead, they have initialized the temperature at the wind base and used this to generate a planetary wind. The recent work of Owen & Adams (2014) moves toward a more self-consistent picture by simulating ionization-driven winds from hot Jupiters, with stellar winds and magnetic fields, in 2D. Still, there is a need for 3D simulations which take into account photoionization, and this motivates our work.

As a first step toward realistic models of hot Jupiter mass loss, we present a new 3D hydrodynamic model of atmospheric escape with self-consistent heating. We focus on how asymmetric heating affects the flow near the planet. The model and results are organized in this paper as follows. In Section 2, we describe the physics included in our simulation. In Section 3 we describe our simulation setup and our initial conditions. In Section 4, we present our results of the time-evolved wind structure, mass-loss rates, and comparisons to 1D models. In Section 5, we provide our estimate of the predicted Lyα transmission spectrum mid-transit. In Section 6, we conclude with a summary and a discussion of future extensions.

2. MODELING HYDRODYNAMIC ESCAPE

To model hydrodynamic escape we conduct 3D radiation hydrodynamic simulations. Although the upper atmosphere of a hot Jupiter is low density, the mean free path remains small compared to the scale height, justifying the fluid approximation. To attain sufficient dynamic range to resolve the upper atmosphere, we include only the planet in our simulation domain. The star resides outside of the computational boundary and exerts its influence through the gravitational potential and through ionizing photons entering a single boundary.

We use the publicly available grid-based code Athena, version 4.1 (Stone et al. 2008), to solve the equations of ideal hydrodynamics. We implement an additional module for photoionizing radiation from Krumholz et al. (2007), as described in Appendix A. Our modified version of Athena, with initial conditions files, is freely available for download and use.[5] In this section, we describe the numerical implementation of our model, and our initial conditions can be found in Section 3.

2.1. Hydrodynamics

We use Athena to solve the following set of hydrodynamic equations, including gravitational, radiative, and chemical evolution source terms:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{v}) = 0, \tag{1}$$

$$\frac{\partial}{\partial t}(\rho \boldsymbol{v}) + \nabla \cdot (\rho \boldsymbol{v}\boldsymbol{v}) + \nabla P = -\rho \nabla \Phi \tag{2}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot ((E + P)\boldsymbol{v}) = \mathcal{G} - \mathcal{L}, \tag{3}$$

$$\frac{\partial \rho_n}{\partial t} + \nabla \cdot (\rho_n \boldsymbol{v}) = \mathcal{R} - \mathcal{I}, \tag{4}$$

where ρ is the total density, ρ_n is the density of neutral gas, $\boldsymbol{v}$ is the velocity, P is the thermal pressure, and Φ is the gravitational potential. The total energy density E, excluding the chemical potential energy, is

$$E \equiv \epsilon + \rho \frac{\boldsymbol{v} \cdot \boldsymbol{v}}{2}, \tag{5}$$

where ϵ is the internal energy density excluding the chemical potential. Omitting the chemical potential allows us to use the usual relationship between pressure and internal energy for an ideal gas,

$$\epsilon = \frac{P}{\gamma - 1}, \tag{6}$$

and to adopt an adiabatic equation of state with $\gamma = 5/3$ as a constant, appropriate for either the atomic or ionized gas expected to be found in the upper atmosphere of a hot Jupiter. Excluding the chemical energy and treating γ as constant in this manner is reasonable because we are in a regime where collisional ionization is negligible, and thus there is a strong separation of scales between the mean particle kinetic and chemical energies. In the energy equation, Equation (3), $\mathcal{G}$ and $\mathcal{L}$ are the rates of radiative heating and cooling, respectively. In the continuity equation for the neutral density, Equation (4), $\mathcal{R}$ and $\mathcal{I}$ are the rates of recombination and ionization, respectively.

2.2. Gravitational Potential

We employ a static gravitational potential from both the planet and the host star, including the contribution from the centrifugal or tidal term:

$$\Phi = -\frac{GM_{\rm p}}{r} - \frac{GM_\star}{r_\star} - \frac{1}{2}\frac{GM_\star r_\star^2}{a^3}, \tag{7}$$

[5] https://github.com/tripathi/Atmospheric-Athena

where M_p is the mass of the planet, $M_\star$ is the mass of the star, a is the distance between the centers of the planet and star, r is the local distance to the planet, and $r_\star$ is the local distance to the star.

2.3. Ionization Balance

As noted in Section 2.1, we track both the neutral and the ionized gas in the computational domain. We model the changes in neutral density based on rates of photoionization and recombination. The former depends on the available stellar ionizing flux, while the latter depends on the number densities of ions and neutrals in the gas.

We consider a simulation box whose distance from the star is significantly larger than the size of the box, $l_{\rm box}$. We neglect geometric spreading of the stellar radiation field and simply treat it as a planar front of radiation entering the box. Since the near edge of the box is 4.5 $l_{\rm box}$ from the star and the far edge is 5.5 $l_{\rm box}$, the change in flux between the near and far edges of the box and the center, where we have normalized, is 20%. We are therefore making an error of this order by neglecting geometric spreading. A planar front of radiation enters the box with photon flux F_0. In this approximation, the flux at a distance l into the box is given by

$$F(l) = F_0 e^{-\tau(l)}, \tag{8}$$

which depends only on the optical depth τ, defined as

$$\tau(l) = \int_0^l n_{\rm H}\sigma_{\rm ph}\, d\ell, \tag{9}$$

where $\sigma_{\rm ph}$ is the photoionization cross-section and $n_{\rm H}$ is the neutral number density.

Given the abundance of hydrogen in hot Jupiter atmospheres, we consider a pure hydrogen atmosphere so that $n_{\rm H} = \rho_n/\mu_{\rm H}$, where the mean gas mass per hydrogen nucleus is $\mu_{\rm H} = m_{\rm H} = 1.67 \times 10^{-24}$ g and $m_{\rm H}$ is the mass of a hydrogen atom. For hydrogen photoionization at the threshold energy of 13.6 eV, the cross-section is $\sigma_{\rm ph} = 6.3 \times 10^{-18}\,{\rm cm}^2$. In reality, the cross-section will depend on the photon frequency as roughly ν^{-3}; however, since the chemical state is mostly controlled by photons near the ionization threshold for optically thin gas and the energetics are controlled by photons that span much less than a factor of 2 in frequency, we simply adopt the threshold cross-section for all purposes. While Equation (9) is true for any frequency, we assume that our ionizing source is monochromatic, for comparison with previous work. We review our choice of monochromatic flux in Section 6. The calculated flux is used to determine the time rate of change of the ionized density, in units of g cm^{-3} s^{-1}, as

$$\mathcal{I}_{\rm ph} = \sigma_{\rm ph}\rho_n F(l). \tag{10}$$

The photoionization rate is simply $\mathcal{I}_{\rm ph}/\mu_{\rm H}$.

Our treatment of photoionization and recombination assumes the case-B condition and the on-the-spot approximation, so that the gas is optically thick to ionizing photons in regions where we consider recombination. For recombinations to the ground level, we assume that emitted ionizing photons will ionize at the same location, effectively canceling out that recombination. Thus, the time rate of change of the recombined density, in units of g cm^{-3} s^{-1}, is given by

$$\mathcal{R} = \mu_{\rm H}\alpha^{(B)} n_{\rm e} n_{\rm H^+}, \tag{11}$$

where $\alpha^{(B)} = 2.59 \times 10^{-13}(T/10^4\ {\rm K})^{-0.7}\,{\rm cm}^3\,{\rm s}^{-1}$ is the case-B recombination coefficient (Osterbrock 1989), $n_{\rm H^+}$ is the number density of protons, and $n_{\rm e}$ is the number density of electrons. We assume that $n_e = n_{\rm H^+} = (\rho - \rho_n)/\mu_{\rm H}$. Note that the recombination rate is simply $\mathcal{R}/\mu_{\rm H}$.

We omit collisional ionization from our treatment because it is negligible for our problem, since MC09 found that temperatures at the atmosphere's wind base are $\sim 10^4$ K. The gas temperature T is determined from the total energy density E and momentum density $\boldsymbol{p}$ using

$$T = \frac{\gamma - 1}{\rho}\left(E - \frac{1}{2}\frac{|\boldsymbol{p}|^2}{\rho}\right)\frac{\mu}{k}, \tag{12}$$

where the mean gas mass, $\mu = x\mu_{\rm i} + (1 - x)\mu_{\rm N}$, depends on the ionization fraction, $x = 1 - \rho_n/\rho$, the mean particle mass of ionized species $\mu_{\rm i}$, and the mean particle mass of neutral gas $\mu_{\rm N}$. For our hydrogen gas, we use $\mu_{\rm i} = m_{\rm H}/2$ and $\mu_{\rm N} = m_{\rm H} = 1.67 \times 10^{-24}$ g. We note that though the true value of μ is $\mu = m_{\rm H}/(1 + x)$ for atomic and ionized hydrogen, the above expression, chosen for compatibility with previous code development, provides a good approximation.

2.4. Heating and Cooling

Radiation not only ionizes the gas, but it also contributes to its heating and cooling. Each ionizing photon imparts its energy in excess of the ionization threshold to the newly liberated electron, which then heats the gas through collisions. Each photon contributes an energy, e_Γ, to heating the gas, so that the photoionization heating rate, per unit volume, is

$$\mathcal{G}_{\rm ph} = e_\Gamma \sigma_{\rm ph} n_{\rm H} F(l). \tag{13}$$

We use an ISM appropriate heating rate of $e_\Gamma = 2.4$ eV (Whalen et al. 2004), chosen to validate our ionizing radiative transfer code against H II region ionization fronts, discussed in Appendix A.2. While we use this heating rate for both our validation tests and our mass loss model, we note that for a hot Jupiter atmosphere absorbing a quiet solar Lyman continuum spectrum, Trammell et al. (2011) calculated $e_\Gamma = 2.7$ eV, assuming 100% energy deposition efficiency. Efficiency values may be smaller (Koskinen et al. 2014), making our choice of heating rate an upper limit for mass loss.

Hydrogen recombination emits photons, which can escape and thus cool the gas. The radiative recombination cooling rate, from a linear fit to values in Osterbrock (1989), is

$$\mathcal{L}_{\rm rec} \approx \left(6.11 \times 10^{-10}\,{\rm cm}^3\,{\rm s}^{-1}\right) kT \left(\frac{T}{\rm K}\right)^{-0.89} n_{\rm e} n_{\rm H^+}, \tag{14}$$

where k is the Boltzmann constant.

We also include the rate for cooling from neutral atoms that are collisionally excited and emit Lyα photons, given by Black (1981) as

$$\mathcal{L}_{{\rm Ly}\alpha} = \left(7.5 \times 10^{-19}\,{\rm erg\,cm}^{-3}\,{\rm s}^{-1}\right) e^{-118348\ {\rm K}/T} n_{\rm e} n_{\rm H}. \tag{15}$$

Menager et al. (2013) suggest that this formula may overestimate the Lyα cooling rate by up to an order of magnitude for hot Jupiter atmospheres, due to non-LTE and other radiative transfer effects. As we discuss in Section 4.1.2, Lyα cooling is negligible almost everywhere for our planet's parameters, so

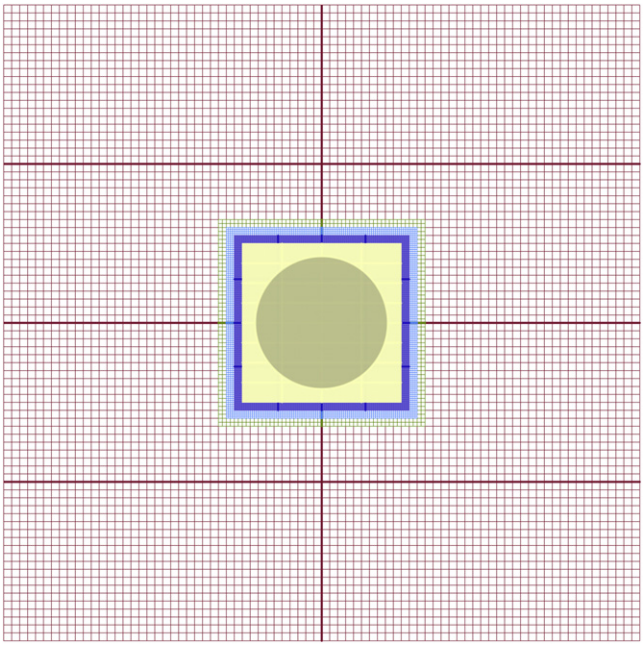

Figure 1. Our computational grid has five levels of resolution (each level is in a different color), to resolve the scale height of the planet (gray). The grid spans 1.5×10^{11} cm ($10R_p$) in each dimension, with a highest resolution of $1/128\ R_p$. This 2D slice in the x–y plane shows how the grid is distributed across processors, as denoted by the thick lines.

such a difference does not change our results, but it may affect hotter and denser planets.

2.5. Numerical Algorithm

We use Athena's Roe's linearized Riemann solver, with default second-order spatial reconstruction of the fluid variables and the directionally unsplit corner transport upwind integrator in 3D (Stone et al. 2008). To avoid the carbuncle instability (Quirk 1994), which we find when photoionization heated gas advects around the planet and converges on the nightside, we use Athena's built-in H-correction (Stone et al. 2008), to add dissipation when strong shocks are aligned with the grid. Our fiducial simulation length is roughly four orbital periods, by which point the wind has reached a steady state.

3. INITIAL CONDITIONS

In this section, we describe our choice of initial conditions for the computational domain, planetary atmosphere, and the host star. As shown in Figure 1, we initialize our computational domain with a planet at the origin of a 3D Cartesian grid. The computational volume corresponds to a physical size of $(10R_p)^3$, discretized into a base grid of 80^3 cells. To resolve the minimum atmosphere scale height when the wind is launched, we use five levels of grid refinement around the planet, yielding a finest resolution of $(1/128\ R_p)^3$. We show in Appendix B that our resolution is sufficient for our results to reach numerical convergence. At this resolution, the influence of the Cartesian discretization on the spherical atmosphere is minimal.

We work in the frame corotating with the planet's orbit. We assume that the planet is at a fixed distance from the star. In this work, we omit the Coriolis force. We do not include the rotation of the planet since hot Jupiters are tidally locked. We inject stellar UV flux into one face $(-x)$ of our box and calculate the optical depth to photoionization from this boundary. For the fluid variables, all faces of our box have outflow boundary conditions.

3.1. Planetary Atmosphere

To study hydrodynamic escape, we retain a large, well-resolved atmosphere to serve as a mass reservoir from which a wind can be launched. The atmosphere is in hydrostatic equilibrium, with a gas density profile

$$\frac{dP}{dr} = -\frac{GM_p\rho}{r^2}, \tag{16}$$

and in agreement with our adiabatic equation of state, we assume a polytrope

$$P = K\rho^{\gamma}, \tag{17}$$

where K is a constant of proportionality and γ is the adiabatic index. Integrating from a radius R_p (with corresponding density ρ_p) to a radius r yields

$$\rho_{atm}(r) = \left[\frac{\gamma-1}{\gamma}\frac{GM_p}{K}\left(\frac{1}{r}-\frac{1}{R_p}\right)+\rho_p^{\gamma-1}\right]^{1/(\gamma-1)}. \tag{18}$$

For an ideal gas, K can be determined from the isothermal sound speed c_s and ρ using

$$K = \rho^{1-\gamma}c_s^2. \tag{19}$$

We normalize these parameters at R_p, the radius where ionizing photons are absorbed and from which the wind is launched.

Resolving the entire planet is computationally prohibitive and unnecessary, so we create an artificial inner boundary for the atmosphere. We set this inner boundary at $r_b = 0.75R_p$, $\approx$4 scale heights below the wind launching point, which is more than sufficient to maintain a reservoir of atmospheric gas. Interior to r_b, the density and temperature profiles are fixed at their hydrostatic values at every timestep, but non-zero fluxes are permitted across the boundary. To prevent the density from diverging at the origin, as expected from Equation (18), we set $\rho(r \leqslant r_0) = \rho_{atm}(r_0)$ at a radius $r_0 = 0.5R_p$. Our model is insensitive to the choice of r_0, so long as it is smaller than r_b by the number of cells used for the Riemann solver's reconstruction method.

To maintain the stability of the simulated atmosphere, we match the pressure at the atmosphere's outer edge to that of a stationary and uniform ambient medium. Unlike the pressure, the density of the ambient medium is discontinuous from the planet. To ensure that the background gas does not influence the development of winds from the planet, we require it to be at a low enough density (or equivalently, a high enough temperature) so that the inevitable accretion of material onto the planet has a minimal effect on the atmospheric structure and wind launching. Combining these constraints, we use the following profile to specify the density of the gas in the entire domain,

$$\rho(r) = \begin{cases} \rho_{atm}(r_0) & r < r_0, \\ \rho_{atm}(r) & r_0 \leqslant r \leqslant r_e, \\ \rho_{atm}(r_e)\cdot 10^{-4} & r > r_e, \end{cases} \tag{20}$$

Table 1
Simulation Parameters

Parameter	Value
Planet	
Mass, M_p (g)	10^{30}
UV photon absorption radius, R_p (cm)	1.5×10^{10}
Density at R_p, ρ_p (g cm^{-3})	10^{-15}
Isothermal sound speed at R_p, $c_{s,p}$ (cm s^{-1})	3×10^5
Hydrogen photoionization cross-section, σ_{ph} (cm^2)	6.3×10^{-18}
Mean mass per hydrogen nucleus, μ_H (g)	1.67×10^{-24}
Star	
Mass $M_\star$ (g)	2.0×10^{33}
Orbital distance, a (cm)	7.48×10^{11}

where r_e is the edge of the planet. The pressure profile is

$$P(r) = \begin{cases} K\rho_{atm}(r_0) & r < r_0, \\ K\rho_{atm}(r)^\gamma & r_0 \leqslant r \leqslant r_e, \\ K\rho_{atm}(r_e)^\gamma & r > r_e. \end{cases} \quad (21)$$

Since we model a spherical planet on a Cartesian grid, the pressure is not perfectly matched between the ambient gas and the edge of the planet's atmosphere. Nevertheless, we find that the changes in our atmosphere over a sound crossing time are substantially smaller than the outflows that develop.

Our fiducial parameters are summarized in Table 1. We model a low mass and extended hot Jupiter with $R_p = 1.5 \times 10^{10}$ cm $= 2.14 R_{Jup}$ and $M_p = 10^{30}$g $= 0.53 M_{Jup}$. These values are similar to WASP-17b, the most inflated exoplanet to-date, with $R_p = 1.97 \pm 0.06 R_{Jup}$ (Bento et al. 2014) and $M_p = 0.477 \pm 0.033 M_{Jup}$ (Southworth et al. 2012). It is important to note that the radius that we quote as R_p for our model corresponds to the absorption radius for ionizing photons, which is larger than the observed optical planetary radius. Our choice of a low mass, low density planet reduces the computational cost to resolve the upper atmosphere and launch an outflow.

Following the parameter study of MC09, which found that mass loss rates are insensitive to wind-base temperatures $\lesssim 10^3$ K, we set $c_{s,p} = 3 \times 10^5$ cm s^{-1}, or equivalently $T_p = 1.1 \times 10^3$ K. Our definition of R_p requires that over a scale height evaluated at R_p, $H = c_s^2/(GM_p/R_p^2)$, the optical depth reaches unity. This corresponds to a number density, $n_p = 6 \times 10^8$ cm^{-3}. To ensure that the gas is optically thick, we extend the planet's atmosphere to a radius $r_e = 1.02 R_p$, where $\rho(r_e) = \rho_p/10$. We begin with a neutral hydrogen atmosphere with the density at R_p, $\rho_p = \mu_H n_p = 10^{-15}$ g cm^{-3}. We note that though we set our initial conditions so that initially $\tau \approx 1$ at R_p, the gas is allowed to self-consistently choose where the $\tau = 1$ surface lies. While the planet begins completely neutral, the background gas is fully ionized. Starting with a fully ionized and optically thin background allows us to track the evolution of the planet alone.

3.2. Stellar Radiation

We include the host star's gravity in our static gravitational potential. We set the mass and radius of the star as $M_\star = M_\odot$ and $R_\star = R_\odot$. The stellar radius is used to calculate the predicted Lyα extinction during transit. Typical of hot Jupiters, the star is located at a distance of 0.05 AU from the planet.

Since hot Jupiters are tidally locked to their host stars, stellar flux is constantly received by the same face of the planet. We include the stellar flux as a plane-parallel source of ionizing radiation incident on one side of our computational box. As described in Section 2.3, the plane-parallel approximation is justified by the relative sizes and orbital separation of the planet and the star. We treat the stellar flux as a monochromatic source, rather than a full stellar spectrum. This choice allows us to make direct comparisons with the MC09 1D model, which also uses a monochromatic source.

We illuminate the planet gradually, slowly increasing the flux over two orders of magnitude, using the function,

$$\frac{F(t)}{F_0} = 5 \ \mathrm{erf}\left(\frac{t}{8 \times 10^4 \ \mathrm{s}} - 1.5\right) + 5.1, \quad (22)$$

for a given choice of F_0. The "ramp up" timescale is chosen to be significantly longer than the advection time for gas moving around the planet. By increasing the photon flux slowly, we allow the system to gradually relax to equilibrium.

To examine changes in mass loss as a function of flux, we run our simulation with two different stellar flux values. For our fiducial model, we use a maximum steady state photon flux of 2.02×10^{14} cm^{-2} s^{-1}, which corresponds to $10.1 F_0$ for $F_0 = 2 \times 10^{13}$ cm^{-2} s^{-1}. Since this value of F_0 is comparable to the solar UV flux from Woods et al. (1998) scaled to our orbital distance of 0.05 AU, our fiducial model represents a relatively younger, hotter star than the Sun. To study a Sun-like star, we also run a second model, where $F_0 = 2 \times 10^{12}$ cm^{-2} s^{-1}. Because we use single frequency photons, there is not a one-to-one correspondence between our choice of F_0 and the solar flux; a different frequency range would necessitate a different F_0.

4. RESULTS

The structure of the planetary outflow is described in Section 4.1, with winds on the day and night sides described in 4.1.1 and 4.1.2, respectively. The mass loss rate for our fiducial model, as well as the lower-flux model, is presented in 4.2. Agreement between our simulation and 1D models is discussed in 4.3.

4.1. Wind Structure

Within an orbital period of illumination, the planet develops a steady-state, transonic atmospheric outflow. After reaching steady state at $\sim 3 \times 10^5$ s, escape continues in steady state for the remainder of our simulation: three orbital periods or four sound-crossing times for our simulation box. The transition to steady-state can be seen in the time evolution of the mass loss rate, discussed in Section 4.2.

The wind's evolution from initialization to steady state is illustrated in Figure 2. This and other midplane visualizations were generated using VisIt (Childs et al. 2012). Since our model is axisymmetric, the outflow structure is the same in both the x–z (shown here) and the x–y planes. The dayside exhibits a strong, ionized outflow. The nightside outflow is slightly weaker, with neutral gas in the planet's shadow. We discuss these day–night differences and the underlying energetic and ionization processes in Sections 4.1.1 and 4.1.2.

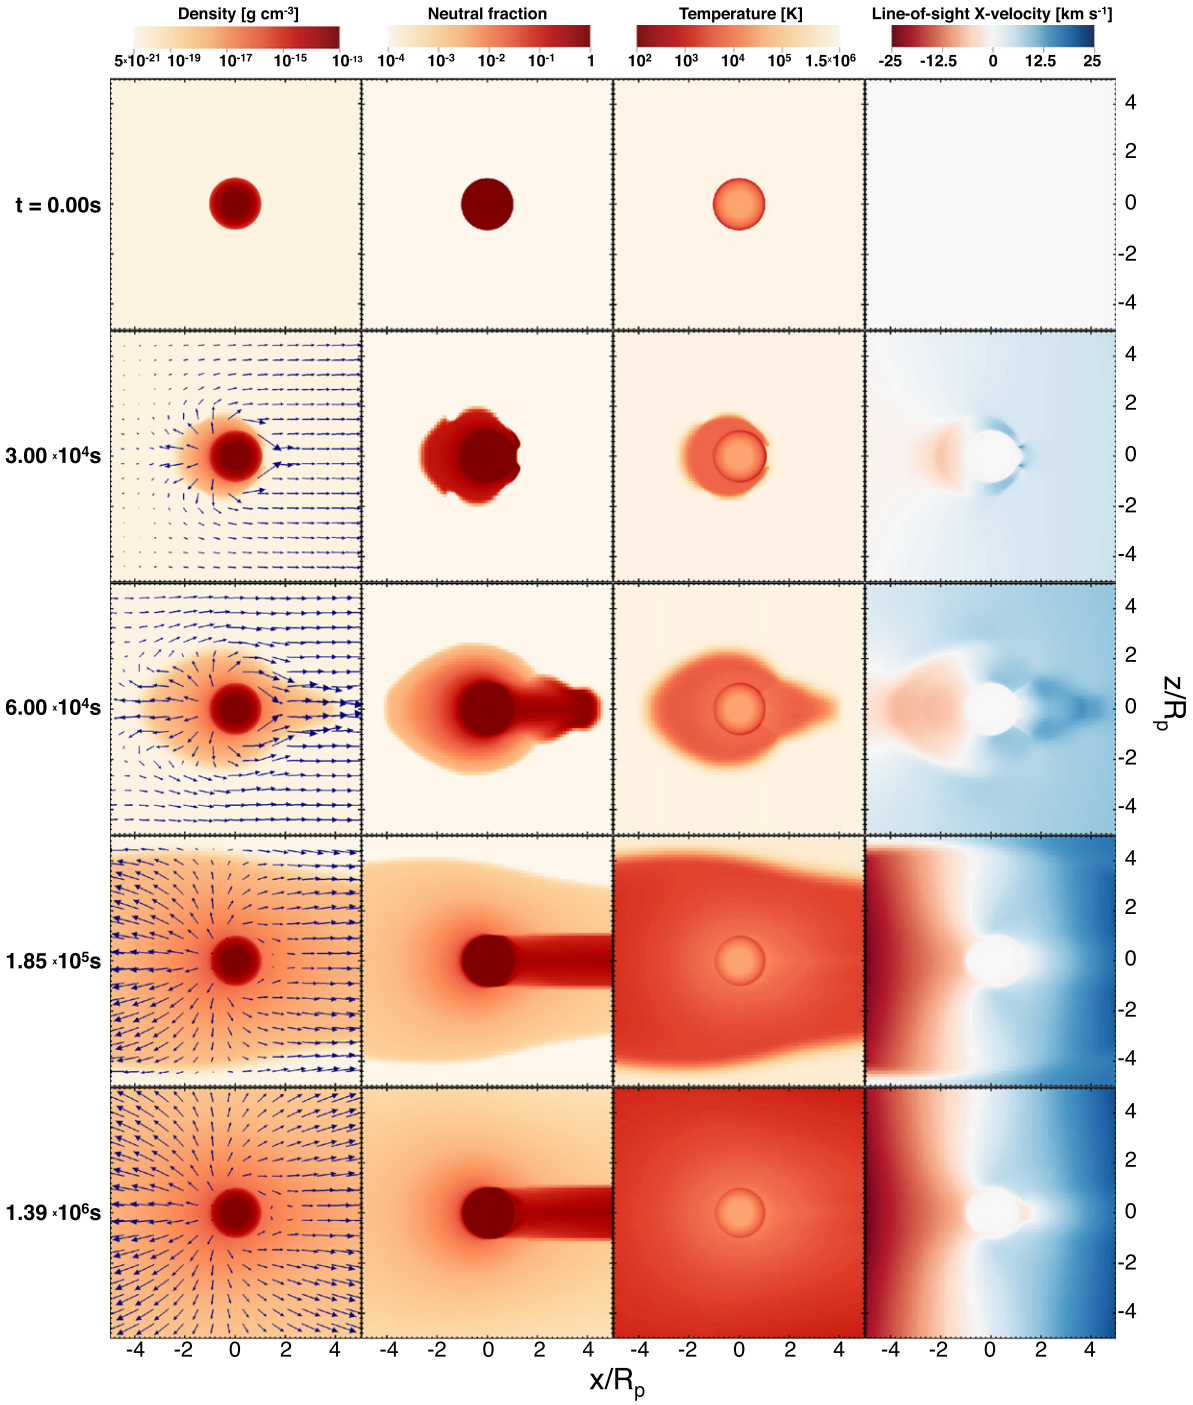

Figure 2. Atmospheric expansion on the dayside, with advection toward the nightside, as seen in this time evolution of density (g cm^{-3}) with velocity vectors overplotted (left), neutral fraction (second column), temperature (K) (third), and x-velocity (km s^{-1}) (right), for a slice through the midplane. Top row ($t = 0$ s): initial conditions, second row (3×10^4 s): the atmosphere heats on the dayside and advects to the nightside, third row (6×10^4 s): the outflows are tidally extended along the axis to the star, fourth row (1.85×10^5 s): the outflows continue at larger radii, with lower density, and finally bottom row (1.39×10^6 s): the steady state outflow. The star illuminates the planet from the $-x$ boundary, and the mid-transit observer is at $+x$.

As shown in Figure 3, the steady state outflow is transonic. Its radial velocity reaches the local escape speed after exceeding the adiabatic sound speed. The flow becomes supersonic at distances from the planet comparable to the Roche lobe radius, $R_{\rm Roche} = [M_{\rm p}/(3M_\star)]^{1/3}a = 4.1 \times 10^{10}$ cm $= 2.7R_{\rm p}$. The escape surface is largely outside of the Roche lobe radius and is closest to the planet along the substellar ray.

Tidal gravity is responsible for elongating the outflow and contributes to the asymmetric escape surface. Consequently, the velocity field of the outflow observed during a planetary transit should be dominated by the line-of-sight velocity.

Throughout the results section, we discuss the wind launching in time, so as to understand its final structure and

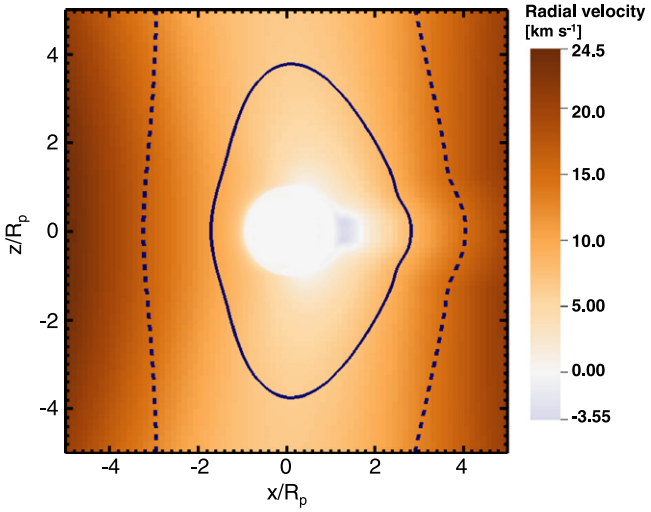

Figure 3. Steady-state outflow radial velocity ($\mathrm{km\,s^{-1}}$), shown with the adiabatic sonic surface (solid) and the escape surface (dashed) in the midplane. For reference, the Roche lobe radius is $2.7R_p$.

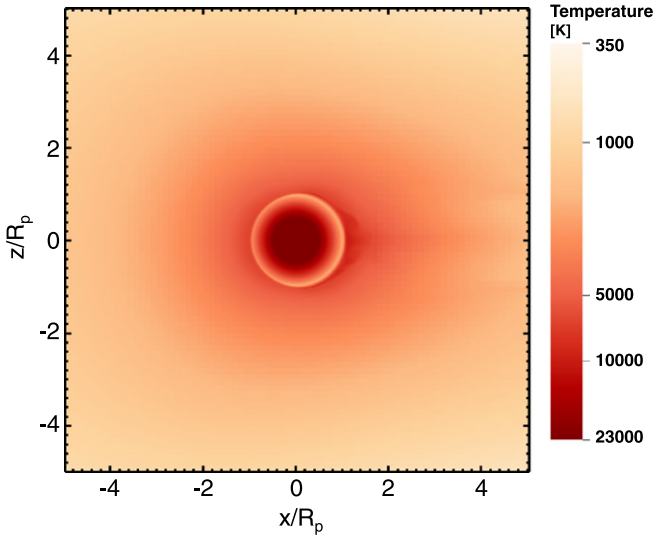

Figure 4. Temperature (K) of the steady-state outflow exhibits a day–night asymmetry, seen in this midplane slice. While dayside temperatures remain $<10^4$ K, neutral nightside gas approaches 10^4 K.

dynamics. Because our initial conditions do not reflect the properties of a newly formed hot Jupiter, this time evolution does not represent the early evolution of a hot Jupiter's outflow. However, the model's time evolution may be relevant for planets orbiting very active, flaring stars, whose flare timescales can be comparable to our simulated flux ramp up time.

4.1.1. Dayside Flow: Directly Launched by Photoionization Heating

On the planet's dayside, photoionization heats the atmosphere to ~7000 K, by depositing energy at R_p, where $\tau = 1$ to photoionization. Irradiated gas above R_p accelerates and moves outwards. This outward expansion, hereafter referred to as PdV work, is the primary coolant of the gas. Radiative recombination cooling is an order of magnitude smaller. Because our planet has low surface gravity, its outflow does not achieve the 10^4 K temperature required for substantial Lyα cooling, as shown in Figure 4. Lower temperatures are sufficient to accelerate gas to the planet's escape speed.

To assess in more detail the relative contributions of various heating, cooling, and ionization processes to the wind's structure, we recast Equations (3) and (4) as

$$-\rho \boldsymbol{v} \cdot \nabla\left[\frac{P}{\rho(\gamma-1)}\right] + \frac{P}{\rho}(\boldsymbol{v} \cdot \nabla\rho) + \mathcal{G}_{\rm ph} - \mathcal{L}_{\rm rec} - \mathcal{L}_{\rm Ly\alpha} = 0 \quad (23)$$

and

$$-\frac{\rho}{\mu}\boldsymbol{v} \cdot \nabla x + \frac{\mathcal{I}}{\mu_{\rm H}} - \frac{\mathcal{R}}{\mu_{\rm H}} = 0. \quad (24)$$

The first two terms in Equation (23) represent the change in internal energy and the PdV work, respectively. The first term in Equation (24) represents the advection of ions out of a given cell.

Above the wind base, PdV cooling and photoionization heating contribute to energy balance, as shown in Figure 5, which displays the terms of Equations (23) and (24) in steady state. Near the wind base, heat is stored in internal energy, as shown by the negative internal energy below $1.1R_p$ along the substellar ray. Farther from the wind base, the outflow is driven by local photoionization heating, as well as the heat that was deposited by photoionization lower in the atmosphere and stored in internal energy.

Due to the large photoionization rate of $\sim 10^4\,\mathrm{cm^{-3}\,s^{-1}}$, the dayside outflow is everywhere ionized. The dayside is in ionization balance above the wind base, as seen in the lower panel of Figure 5. Near the wind-base, photoionization contributes a steady source of ions, which are advected away. Ion advection is only important near the wind base, where the fraction of the gas that is ionized is low enough that recombination is too slow to balance ionization. Above $1.1R_p$, recombination balances photoionization.

4.1.2. Nightside Flow: Advected Ions Mixed with a Neutral Shadow

As seen in Figure 2, heated dayside gas not only moves radially outwards, but also advects around the planet. The advected flow exceeds the sound speed and escape velocity, before leaving the box in a steady-state wind. The nightside flows are aligned along the anti-stellar ray, in part due to tidal gravity.

Unlike the dayside flow, the nightside wind is not directly driven by photoionization heating. Instead, flows moving around the planet converge on the nightside, at a stagnation point, and heat the surrounding gas to $\sim 10^4$ K. Below the stagnation point, PdV work adds to the internal energy. Above the stagnation point, internal energy drives the outflow, which then cools by PdV expansion. The interplay between PdV work and internal energy storage allows the gas to be in an energetic equilibrium. With the exception of localized Lyα cooling near the stagnation point, where the gas is hot enough, other sources of cooling, including recombination, are negligible.

In the planet's shadow, the gas has a sharply defined neutral outflow. As highlighted in Figure 6, the planet's nightside atmosphere contributes neutral gas to this flow, in a time-varying circulation pattern. This gas originates above the atmosphere's initial outer radius of 1.56×10^{10} cm $= 1.04R_p$, rather than the dayside value of R_p, since there is no photoionization in the shadow to move away material below

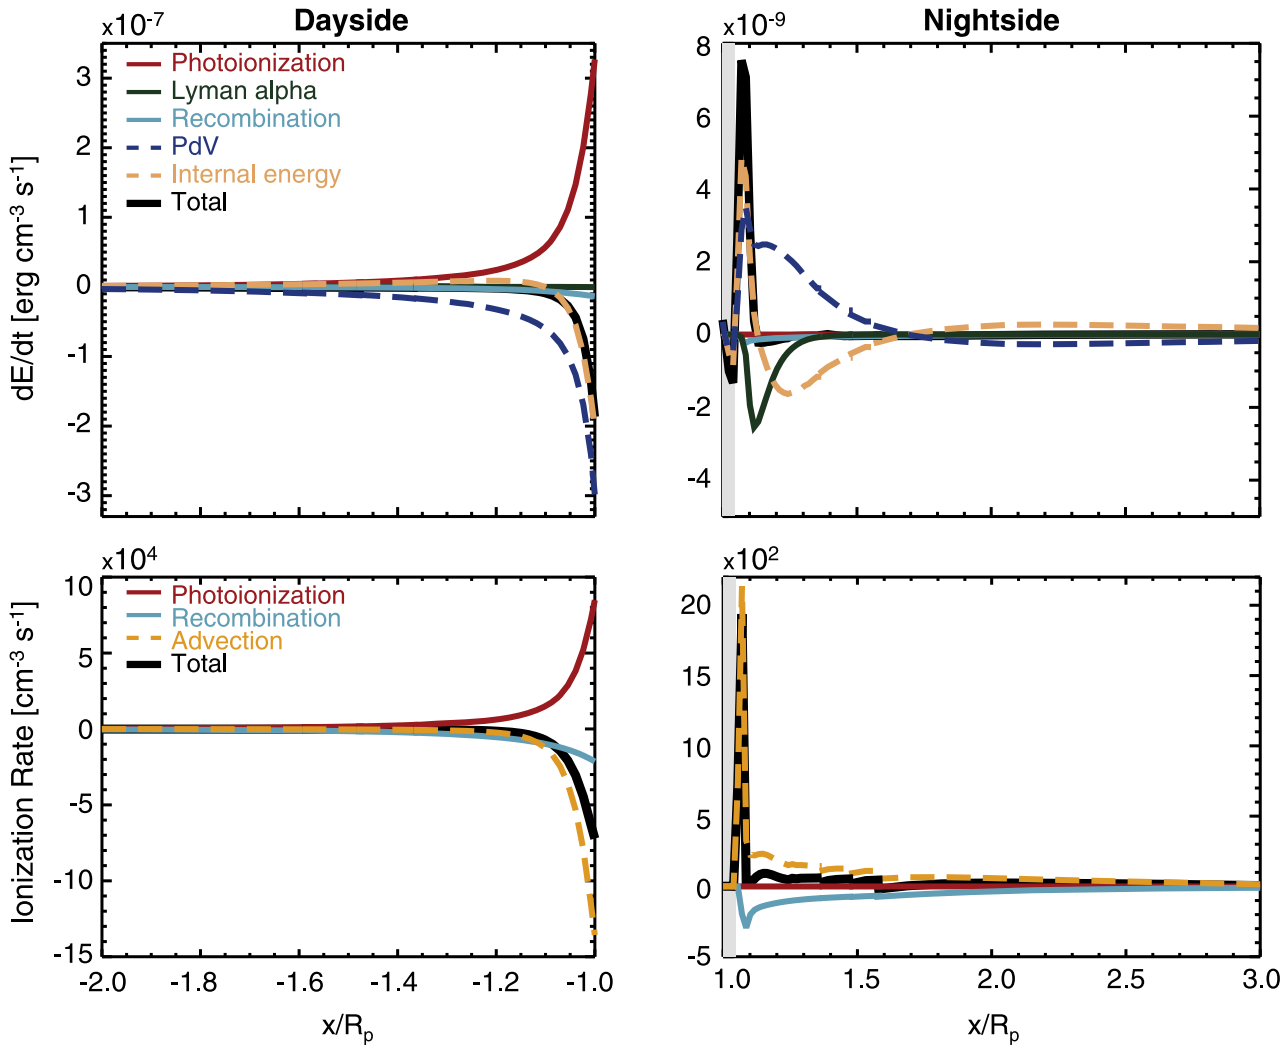

Figure 5. Steady-state energy (upper) and ionization (lower) balance along the substellar ray on the dayside (left) and nightside (right), showing the terms in Equations (23) and (24). Above the wind base on both sides, the escaping gas is in equilibrium, as shown by the total at 0. The shaded region highlights gas below the outflow base, on the nightside.

this radius. We note that this is the only portion of the flow that retains memory of our atmospheric initial conditions. As shown in Figure 4, this nightside neutral gas is also some of the hottest in our computational volume, due to enhanced heating from the converging flows.

The circulation is mostly confined within a stagnation point where the advected, dayside gas converges. Periodically the circulation cell, which may be unresolved turbulence, grows large enough so that the circulating neutral gas mixes with the advected ionized gas and then gets dragged outwards, beyond the stagnation point. Since the planet is optically thick to photoionization, the only ionized gas in this outflow is from advection. Outside of the planet's shadow, the gas is optically thin to photoionization and, thus, it is ionized. Recombination increases the nightside neutral fraction, but only slightly, because the recombination timescale is long compared to the time it takes gas to flow out of our box.

4.2. Mass Loss Rates and Low-flux Results

We calculate the instantaneous mass loss rate by taking the average of the mass flux through a spherical shell, with a finite thickness of one cell, using

$$\langle \dot{M}(r) \rangle = \langle \rho v_r \rangle 4\pi r^2. \quad (25)$$

For comparison, we also calculate the mass loss flux through the faces of a cube with dimensions $2r$ and obtain comparable mass loss rates.

As shown in Figure 7, we find a steady-state mass loss rate of $1.9 \times 10^{11}\,\mathrm{g\,s^{-1}}$ for our fiducial model. It is interesting to note how the mass loss rate tracks the input flux over time. Before the flux plateaus to its constant value, the mass loss rate is larger at smaller radii—indicative of the time it takes for the heated gas to expand and escape. The mass loss rate reaches steady-state at $\sim 3 \times 10^5$ s, just after the flux reaches its constant value of $2.02 \times 10^{14}\,\mathrm{cm^{-2}\,s^{-1}}$. In steady-state, the mass loss rate is the same at all radii above the Roche lobe radius of $2.7R_p$.

To examine the mass loss rate's scaling with UV flux, we also carry out an identical simulation with lower incident flux. For a simulation with a maximum photon flux of $2.02 \times 10^{13}\,\mathrm{cm^{-2}\,s^{-1}}$, an order of magnitude lower than our fiducial model, we find the outflow evolves with a similar

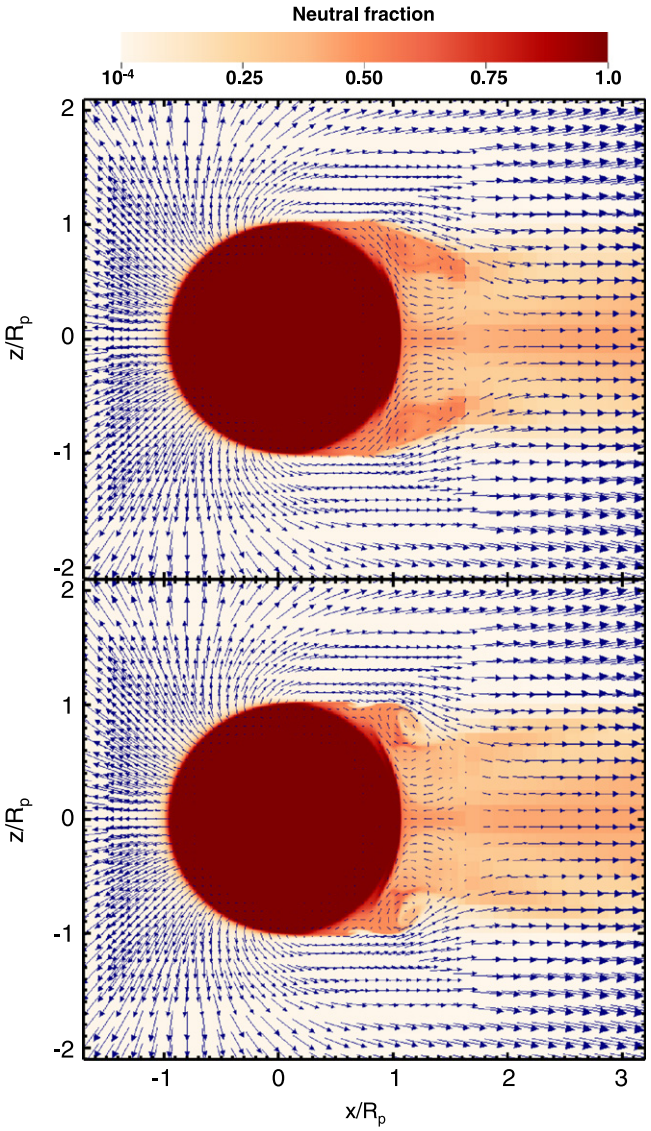

Figure 6. Velocity vectors overplotted on the neutral fraction midplane slice show gas circulating on the nightside of the planet, with ever changing patterns of rolls, shown at 5×10^5 s (upper) and 6×10^5 s (lower). This motion ultimately contributes neutrals to the outflow in the planet's shadow. A stagnation point is visible along the anti-stellar ray at $\sim 1.7 R_p$, between the inward circulating material and the outflowing material.

structure to that of the high flux model, described in Section 4.1. This lower flux simulation reaches a steady state mass loss rate of 2.2×10^{10} g s^{-1}.

4.3. Agreement with 1D Models

We compare a 1D slice from our simulation to the model of MC09 re-calculated for our parameter values. The 1D model of MC09 is a good point of comparison because it also includes ionization balance, tidal gravity, heating and cooling terms. While both the day and night sides are available from our 3D model, only the dayside is available from the 1D models.

Figure 8 shows the steady-state density, velocity, temperature, and ionization profiles from both models, along the substellar ray. We find that the $\tau = 1$ surface, denoted by an x, is at $-R_p$ in both models. Above the $\tau = 1$ surface, the two wind solutions are comparable. Both models transition from

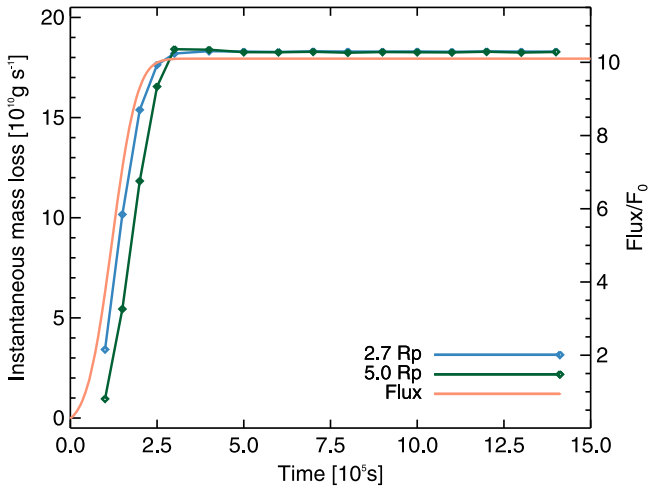

Figure 7. Mass loss rates, calculated at the Roche lobe ($2.7R_p$, blue) and the edge of the box ($5R_p$, green), reach a steady state value of 1.9×10^{11} g s^{-1}, for a model with an incident flux of 2.02×10^{14} cm^{-2} s^{-1}. The evolution of the mass loss rate tracks the incident UV flux (orange).

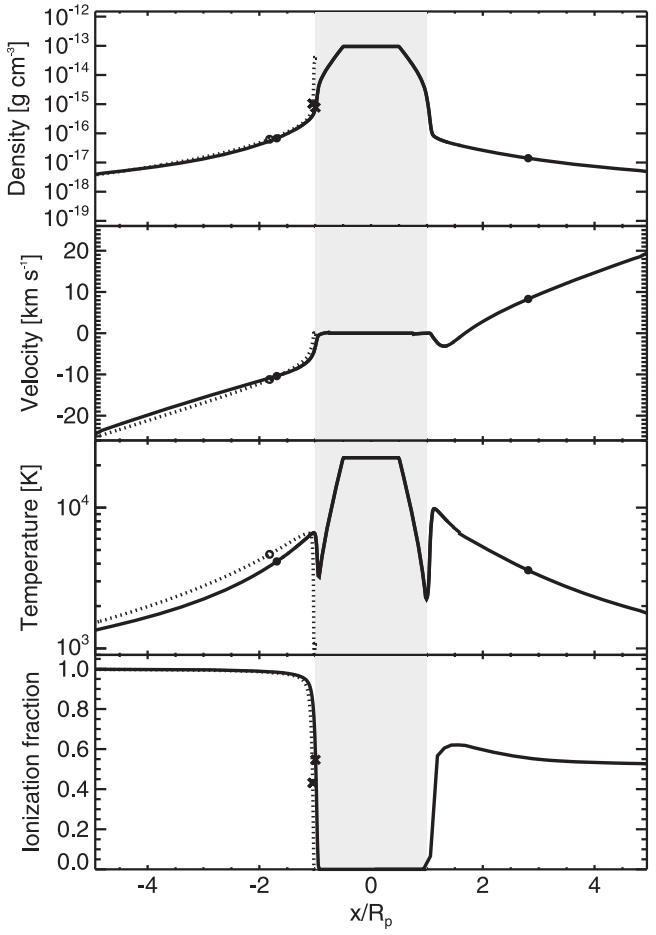

Figure 8. Agreement between our model along the substellar ray (solid) and the 1D model of Murray-Clay et al. (2009) (dotted), for the same parameter values. Circles indicate sonic points and x's denote the $\tau = 1$ surface to photoionization. The unshaded regions are those that we model physically.

subsonic to supersonic velocities at similar dayside locations relative to the planet. We define this sonic point using each model's adiabatic sound speed, rather than isothermal sound speed. The dayside sonic point is interior to the planet's Roche

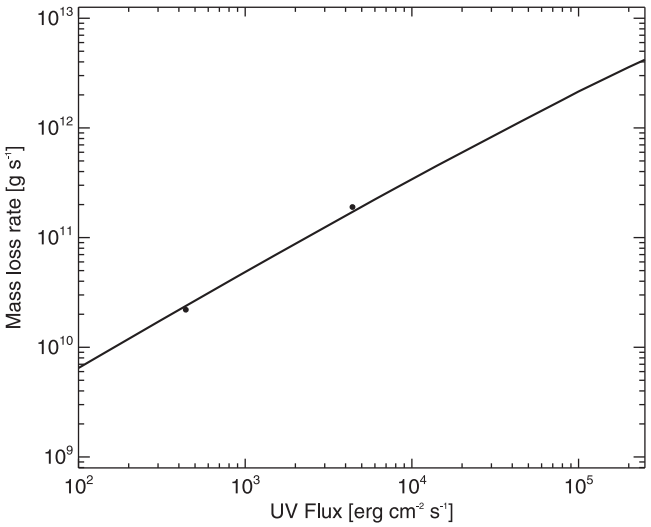

Figure 9. Mass loss rates as a function of UV flux from this work (points) and the 1D model of MC09 (solid line), rerun for the parameters used in this work. The mass loss rate scales as $\dot{M} \propto F_{\rm UV}^{0.9}$ for energy-limited escape.

lobe radius ($2.7R_{\rm p}$). Wind temperatures from both models agree within 25%. The verification of our choice of initial parameters, i.e., $\tau(R_{\rm p}) = 1$, and agreement with comparable 1D results based on the models of MC09 validates that our model is freely and self-consistently setting its parameters, rather than having fixed boundary conditions.

For both models, the $\tau = 1$ location coincides with the sharp transition in ionization fraction at the wind base. In the original MC09 high flux models of a less inflated planet, outflows with such a rapid ionization change are characterized as radiative/recombination limited. However, our highly inflated planet's outflow is energy-limited. The majority of energy deposited by photoionization is converted into PdV work. Advection, rather than recombination, balances photoionization at the very base of the wind. Our planet's highly inflated parameters put its outflow near the boundary of the two regimes.

We compare our mass loss rates to the MC09 1D model, rerun for our stellar and planetary parameters. We estimate the mass loss rate for each 1D model at a radius r by multiplying the substellar mass loss rate by $4\pi r^2$ and a geometric correction factor, to account for differences in the received flux over the planet's surface. For energy-limited outflows, MC09 used a constant correction factor of 0.26, derived from mass loss models with different fluxes that showed $M \propto F_{\rm UV}^{0.9}$.

Our mass loss rates for maximum photon fluxes of 2.02×10^{14} and 2.02×10^{13} cm^{-2} s^{-1}, respectively, agree with the 1D model loss rates and UV flux scaling, as shown in Figure 9. We are unable to make direct comparisons to other mass loss predictions, due to our choice of model parameters. However, given that MC09 had mass loss rates in good agreement with other 1D models for HD 209458 (Yelle 2004; Tian et al. 2005; García Muñoz 2007), our agreement with the MC09 model recalculated for our parameters suggests agreement between ours and other previous 1D models.

5. PREDICTED LYα EXTINCTION

For outflows from our pure hydrogen planet, we calculate the predicted obscuration of stellar Lyα emission during transit.

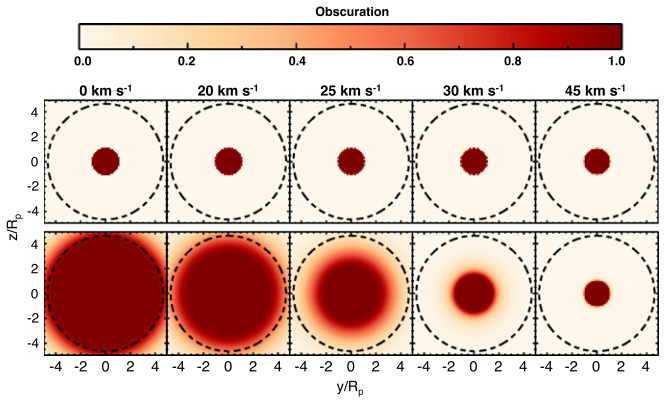

Figure 10. Channel maps of the spatial distribution of Lyα obscuration at line center and then off-center by the specified velocities. The upper panel shows the initial obscuration and the lower panel shows steady-state outflow obscuration, at 1.39×10^6 s. The dashed circle shows the spatial extent of the stellar disk.

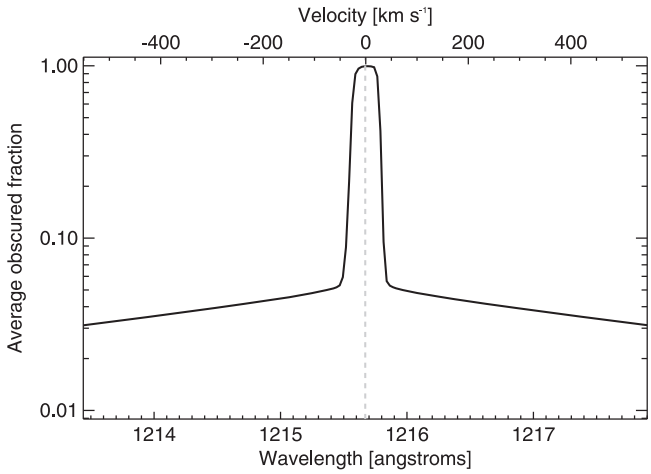

Figure 11. Average Lyα obscuration integrated over the stellar disk, caused by steady-state outflows.

We focus on Lyα because it has been observed for several planets including HD 209458b, HD 189733b, and GJ 436b. Our calculation serves as a first step toward motivating future observations of a wider range of exoplanets, including WASP-17b.

To determine the spatial extent of Lyα absorption, we use the temperature and ionization fraction of each cell in our 3D simulation to compute a Voigt line profile, Doppler shifted by the cell's bulk radial velocity. We use values from Morton (2003) for the Lyα oscillator strength, Einstein A coefficient, and line center wavelength. In this work, we calculate the obscuration produced by a planet passing across the stellar disk mid-transit. We assume that the star has $R_\star = R_\odot$.

Given that our simulation is in 3D, we can examine the integrated absorption spatially, in 2D. Channel maps of the obscuration, $1 - \exp(-\tau)$, initially and in steady-state, are shown in Figure 10. Extinction from the outflow is axisymmetric across our simulated volume. The velocity dependence of the extinction is shown in these maps and more clearly seen in Figure 11, which shows the total obscuration averaged over cells in the stellar disk. To make this figure consistent with traditional sign conventions, the velocity signs have

been reversed and are thus the opposite of other velocity values in this paper. The obscuration is roughly symmetric about line center and drops to 5% at $\pm\approx 50\,\mathrm{km\,s^{-1}}$. At higher velocities, in the line wings, obscuration of a few percent comes only from the planetary disk. Confounding geocoronal emission and absorption by the interstellar medium make observations of Lyα absorbtion challenging near line center, where wind obscuration is most prominent.

6. SUMMARY AND DISCUSSION

We have simulated global hot Jupiter outflows in three-dimensions (3D). The outflow is self-consistently launched by photoionization heating, which we simulate using a new implementation of planar radiative transfer in the Athena hydrodynamic code. Our code is compatible with multiple levels of static mesh refinement (SMR), distributed across parallel processors, and is freely available for further use.

We drive outflows from a highly inflated hydrogen planet, illuminated by a plane-parallel source of ionizing radiation. We find that photoionization heated supersonic outflows emerge on the dayside and advect around the planet, launching nightside outflows not captured in earlier 1D simulations. Outflows are everywhere ionized, except in the planet's nightside shadow, where the outflow is largely neutral. On the planet's night side, a stagnation point in the flow separates outflowing gas from a time-variable circulation region. Perhaps counterintuitively, the longer residence time of gas in this region allows it to reach temperatures higher than those achieved on the planet's dayside. For our fiducial parameters, gas within the nightside stagnation region is the only portion of the flow that is substantially radiatively cooled.

For the highly inflated planet we consider, we find mass loss rates of 1.9×10^{11} and $2.2 \times 10^{10}\,\mathrm{g\,s^{-1}}$ for photon fluxes of 2.02×10^{14} and $2.02 \times 10^{13}\,\mathrm{cm^{-2}\,s^{-1}}$, respectively. The outflow is marginally energy-limited. Our mass loss rates are consistent with the 1D escape model of MC09, rerun for the same parameters. Along the dayside substellar ray, we also find remarkable agreement between the outflow structure in our simulations and the 1D model.

A benefit of our 3D model is the ability to examine not only day and night differences in outflows but also position-dependent extinction. Neutral, shadowed gas, is a major contributor to the absorption of stellar Lyα emission—a key predicted observable. Integrating through the box, we find that the Lyα absorption of the escaping gas obscures stellar emission out to $\approx\pm 50\,\mathrm{km\,s^{-1}}$.

The work presented in this paper is the starting point for more realistic modeling of atmospheric escape. While we have simulated asymmetries in stellar heating, the outflow will also be shaped by the Coriolis force, magnetic fields, and interactions with the stellar wind. In addition, true atmospheres are irradiated by a full spectrum of energetic photons. As our work is motivated by observations of mass loss from close-in planets, in the following paragraphs we review how these physical effects may change the observed transmission spectrum. Since Lyα observations do not yet exist for our planet or the similar WASP-17, we discuss the impact of these physical processes in the context of the observed high velocity Lya obscuration at $\pm 100\,\mathrm{km\,s^{-1}}$ for HD 209458b (Vidal-Madjar et al. 2003). Bear in mind that the two planets have different parameters and thus will not have identical transmission spectra, even with additional physics.

To distinguish 3D geometric effects on the outflow, we have used a monochromatic flux source and compared our results to 1D models, which also used monochromatic flux sources. Energetic effects involved in using a full stellar spectrum were not studied in this first 3D photoionization study. Due to the wavelength dependence of the photoionization cross-section, $\sigma_{\rm ph}$, a full stellar spectrum can smear out the $\tau = 1$ surface where photons are absorbed, increasing the thickness of the wind base (Trammell et al. 2011; Koskinen et al. 2013, 2014). These changes may increase the column density near the base of the wind or increase the radial extent of neutral hydrogen (Koskinen et al. 2013). These effects can contribute to a broadened line profile. At very high incident fluxes, photoionizing X-rays, which have a lower photoionization cross-section per hydrogen nucleus and hence are deposited deeper in the atmosphere, can inject heat quickly enough to be a dominant driver of the outflow (Owen & Jackson 2012). The resulting wind structure is analogous to that modeled for UV-driven winds but begins deeper in the atmosphere, increasing the total wind column. To study this effect, additional spectral bins including photoionization by X-rays can be included in future versions of our model.

High velocity neutral gas may arise in a bow shock that marks the interaction between the outflow and the stellar wind. Tremblin & Chiang (2013) find that charge exchange between outflowing neutral atoms and faster stellar wind ions in this wind-wind interaction region can produce a sufficient fast neutral population.

Trammell et al. (2011, 2014) have suggested that if the planet's magnetic field is dipolar, a magnetically confined dead zone near the equator does not participate in the outflow and could host an enhanced column of neutral gas. MC09 estimate that if the line-of-sight column of neutral gas were enhanced by a factor of 3–5, then naturally broadened line wings from this slower-moving population would be sufficient to produce the observed obscuration (see Figure 11; see also Ben-Jaffel 2008). Regardless of whether magnetic confinement enhances the neutral column, Adams (2011) demonstrates that for reasonable hot Jupiter magnetic field strengths, the magnetic field will redirect and possibly reduce the atmospheric outflow, affecting its 3D structure. For our fiducial model, we find that a local field strength $B \gtrsim 0.06$ G is sufficient for the magnetic pressure $B^2/8\pi$ to overcome the ram pressure of the wind at all radii. For example, at the sonic point of the wind located at $\sim 2R_{\rm p}$, this local B corresponds to that from a dipole field with a surface strength of 0.5 G. Owen & Adams (2014) find using 2D models that increasing magnetic field strength can suppress nightside outflows and reduce planetary mass loss rates. Because we have used Athena, an MHD code, for our simulations, it can be extended to study the effects of the magnetic field on planetary wind confinement.

A line-of-sight column density enhancement could alternatively result from the confinement of the atmospheric outflow by the stellar wind. Depending on the level of stellar activity, pressure confinement by the stellar wind can channel outflowing gas into a smaller solid angle, increasing its local density (Stone & Proga 2009). In extreme cases, the flow is confined to a cylindrical region on the planet's night side. The

Coriolis force (due to the planet's orbital motion) will limit the extent of this high-density region by causing the outflow to curve. The Coriolis force bends trajectories on a timescale $\sim\Omega^{-1}$, where Ω is the angular frequency of the planet's orbit. For our outflow velocity of $v \sim 20\,\mathrm{km\,s^{-1}}$, the length scale of this curvature is $L_{\mathrm{curve}} \sim v\Omega^{-1} \sim 10^{11}\,\mathrm{cm}$, about twice the Roche lobe radius.

Finally, we note that the Coriolis force will produce curvature in the neutral shadow highlighted here and in the 2D model of Owen & Adams (2014). Without such curvature, this large reservoir of neutral gas does not contribute to the line of sight absorption signature since the planet lies between it and the star. At distances larger than L_{curve}, the neutral gas will curve out of the planet's shadow, where it has the potential to substantially enhance the transit absorption signature on the side of the planet opposite to its direction of motion. Upon leaving the nightside, a neutral atom will ionize within $\sim(F\sigma_{\mathrm{ph}})^{-1}\sim$ a few hours, sufficient time to travel approximately a planetary radius at $20\,\mathrm{km\,s^{-1}}$.

Incorporating the Coriolis force into our simulation will clarify how stellar photoionization changes the outflow structure and allow us to produce more realistic, spatially resolved extinction maps and spectra. With the introduction of a stellar wind and magnetic fields, comparisons can be made to 3D studies of these effects that did not include self-consistent photoionization driven wind launching (Cohen et al. 2011; Trammell et al. 2011; Matsakos et al. 2015).

While the model itself can be expanded to include additional physics, post-processing will also offer avenues for other comparisons to observation. Our hydrogen results can be translated into other species, which can be used to look for observational signatures of outflows. Given the 3D nature of our code and the ability to track the time evolution of the gas, we can use a larger computational box to look for sightline specific effects that would appear at ingress and egress. Such a model would allow us to make comparisons to observations with sufficient time resolution, including Kulow et al. (2014) and Lecavelier des Etangs et al. (2012). With improved physics and simulated observational output, we will be one step closer to motivating and understanding observations of mass loss from hot Jupiters.

We thank Xuening Bai, Jim Stone, and James Owen for helpful discussions. We are grateful to the referee for helpful comments that improved the paper. A.T. was supported by the National Science Foundation Graduate Research Fellowship under grant No. DGE-1144152. K.M.K. was supported in part by NASA through Hubble Fellowship grant no. HF-51306.01 awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS 5-26555. R.A.M. acknowledges support from NSF grant AST-1411536. M.R.K. acknowledges support from NSF grant AST-0955300. The computations in this paper were run on the Smithsonian Institution High Performance Cluster (SI/HPC) and the Odyssey cluster supported by the FAS Science Division Research Computing Group at Harvard University.

APPENDIX A
IONIZING RADIATIVE TRANSFER SETUP AND VALIDATION

A.1. Implementation with SMR and Parallelization

We have extended the ray-tracing radiative transfer method from Krumholz et al. (2007) to Athena grids with multiple levels of SMR, parallelized with MPI. We use a plane-parallel source of ionizing radiation, with rays aligned with the grid. As a result, we loop over cells along the direction of radiation propagation and iteratively carry out radiation updates. Radiative source terms, given on the right hand side of Equations (3) and (4), are calculated and used to update the neutral fraction and energy density, using the operator-split approach described in Krumholz et al. (2007). Key to calculating the photoionization source terms is the ionizing flux in each cell. For rays passing through different levels of resolution, we communicate ionizing fluxes across each level and processor boundary.

For a given timestep, our steps for ionizing radiative transfer on a grid with multiple SMR levels, from coarsest (level 0) to finest resolution are:

1. Integrate the flux along rays on level 0.
2. Compute ionization, heating, and cooling rates. Update the energy density and neutral fraction in each cell.
3. Evaluate stopping criteria on level 0, and store the iteration timestep.
4. Prolongate the flux from the current, coarse level (e.g., level 0) onto the next level of resolution (e.g., level 1), where the coarse grid first crosses into the finer grid. To prolongate the flux, copy the coarse flux into the overlapping fine cells, so that each coarse level ray subdivides into eight rays with equal flux, at the next level of resolution.
5. Integrate the flux along the next level of resolution (e.g., level 1), and repeat step 2.
6. Continue this level's integration and radiation update until the coarsest, level 0, timestep has been reached.
7. Repeat steps 4–6 for the remaining resolution levels.
8. Restrict the neutral fraction and energy density from finer, higher resolution levels to coarser levels, e.g., level 2 to level 1, level 1 to level 0, by averaging the finer values of the eight overlapping cells to obtain one value for the coarse level ray.

As described above, when SMR is used, the coarsest grid level is evaluated first, followed by subsequently finer levels. For each level, we integrate the flux at each cell using Equation (8) with a discrete calculation of the optical depth, Equation (9). Calculated fluxes are used to compute radiative source terms and update the neutral fraction and energy density in each cell. As these radiation updates are computationally cheap relative to hydrodynamic updates, multiple radiation updates can be carried out before entering the hydrodynamics integrator. We use the stopping criteria of Krumholz et al. (2007) to determine when the coarse grid radiation update should terminate. To ensure that propagation speeds are the same across different resolution levels, we require finer levels to complete radiation updates for the same elapsed time as that of the coarse level.

With SMR, the flux is needed at the boundaries between different levels of resolution, so we prolongate the radiative flux from lower resolution levels to cells in overlapping higher level cells. Resolution level boundaries and overlapping regions are determined during the simulation initialization. Prolongation occurs at every time step, only at level boundaries, since the flux is then iteratively calculated along grid-aligned rays. Photon number is conserved across level boundaries, since we copy the radiative flux in each coarse cell into the finer cells, which have smaller areas. Recall that Athena requires each level of resolution to be smaller by a factor of 2, in each dimension.

Following the radiation update, we do not restrict the radiative flux from the fine grids back to the coarse grids. In our mass loss simulation, the ambient gas is optically thin and the planet is optically thick, so any flux that needs attenuating will be visible on the coarse levels as needed. We do, however, include a restriction step from the finest levels back to the coarse levels, to recombine finer rays together, for the hydrodynamic variables, E and ρ_n.

When using parallelization, extra communication is required to coordinate, send, and receive ionizing fluxes on different processors. SMR prolongation and restriction operations are complicated by the fact that values are distributed on different processors. To appropriately direct fluxes to the grid levels spread across multiple processors, we store and communicate processor identifiers for the underlying coarse and overlapping fine grids. With this information, flux prolongation now occurs as two separate steps—a sending of the flux values from the coarse grid and a receipt of the values on the fine grid.

Although parallelization speeds up our simulation, there are restrictions. For our plane parallel attenuation of the flux, grids further downstream in the direction of radiation propagation must wait for integration to happen upstream. While radiation updates are computed on the grids in order along the direction of radiation propagation, there are no barriers to carrying out radiation updates on the grids distributed orthogonally to the radiation propagation. As previously noted, we also require coarse levels to complete their radiation updates before finer, overlapping regions can.

A.2. Implementation Validation

We examine the accuracy of our radiation module with SMR and parallelization, by comparing our code to analytic expectations for a planar ionizing source incident on a uniform, optically thick neutral medium. The evolution of the ionized region is akin to that of an H II region Strömgren sphere, but with a planar geometry.

We set up a $(50\,\mathrm{pc})^3$ box of neutral gas ionized by a planar source with a photon flux of $3.67 \times 10^9\,\mathrm{cm^{-2}\,s^{-1}}$, as in Gendelev & Krumholz (2012). To test our implementation, we use a grid with four levels of resolution, distributed among 8 processors, with each processor hosting four grids—one for each level of resolution. All resolution levels were centered within the box, with the coarsest three levels each having 32 cells, and the finest resolution level having 48 cells. To appropriately make comparisons with an H II region, we initialize the gas to an ISM appropriate chemistry. Considering atomic hydrogen with helium, we use the mean gas mass per hydrogen nucleus of $\mu_\mathrm{H} = 2.3 \times 10^{-24}\,\mathrm{g}$ and the mean mass, $\mu = 2.1 \times 10^{-24}\,\mathrm{g}$. We use a carbon abundance of

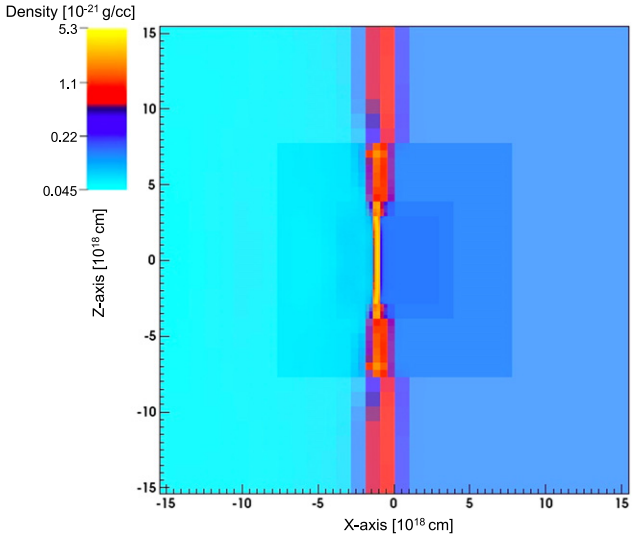

Figure 12. 2D midplane slice showing the density of an ionization front in a uniformly dense gas, simulated with four levels of resolution. To highlight resolution level boundaries, densities in this slice have been rendered with different levels of transparency for each resolution level—coarse levels are more transparent. Radiation is incident from the $-x$ direction.

$\alpha_\mathrm{C} = 3 \times 10^{-3}$ (Sofia & Meyer 2001). We initialize the gas to a uniform neutral medium with number density $n_0 = 63\,\mathrm{cm^{-3}}$ and an isothermal sound speed $c_{s,0} = 5.4 \times 10^4\,\mathrm{cm\,s^{-1}}$. Unlike our hydrogen mass loss simulation, we use $n_\mathrm{e} = n_{\mathrm{H}^+} + \rho\alpha_\mathrm{C}/(14\mu_\mathrm{H})$, where α_C is the carbon abundance of the gas. Since neutral carbon has an ionization threshold below 13.6 eV, C^+ is the dominant ionization state in the neutral ISM, and this provides a minimum abundance of free electrons.

For H II regions, different radiative source terms are needed in the fluid equations. We include all terms from Sections 2.3 and 2.4 except for Lyα line cooling. Additionally, we add contributions from collisional ionization to our ionization balance, using the rate per unit volume of Tenorio-Tagle et al. (1986)

$$\mathcal{I}_\mathrm{coll} \approx (5.84 \times 10^{-11}\,\mathrm{cm^3\,s^{-1}}) n_\mathrm{e}\rho_n \sqrt{\frac{T}{\mathrm{K}}}\, e^{(-13.6\,\mathrm{eV})/kT}. \quad (26)$$

We also include the optically thin cooling terms for molecular cooling, line cooling from O, N, and Ne, and free–free cooling, as described in Krumholz et al. (2007). A full inclusion of heating and cooling terms is necessary to thermostat the ionized gas to 10^4 K.

We can qualitatively assess the success of our radiative transfer implementation by looking at a 2D slice of the gas density, shown in Figure 12. Note that to make the different SMR regions visible, we rendered the product of the density and a constant transparency value, with greater transparency values for coarse resolutions. Radiation incident from the left edge of the box photoionizes the gas, increasing its temperature and lowering its density. The ionized region is separated from the neutral region by an ionization front (red), which we resolve to one grid cell. As a result, the front appears thinner in the higher resolution regions, and it is also denser—to conserve mass. The ionization structure is more clearly seen in Figure 13, which gives a 1D view of the density, temperature, neutral density, and ionization fraction. The ionization front is visible

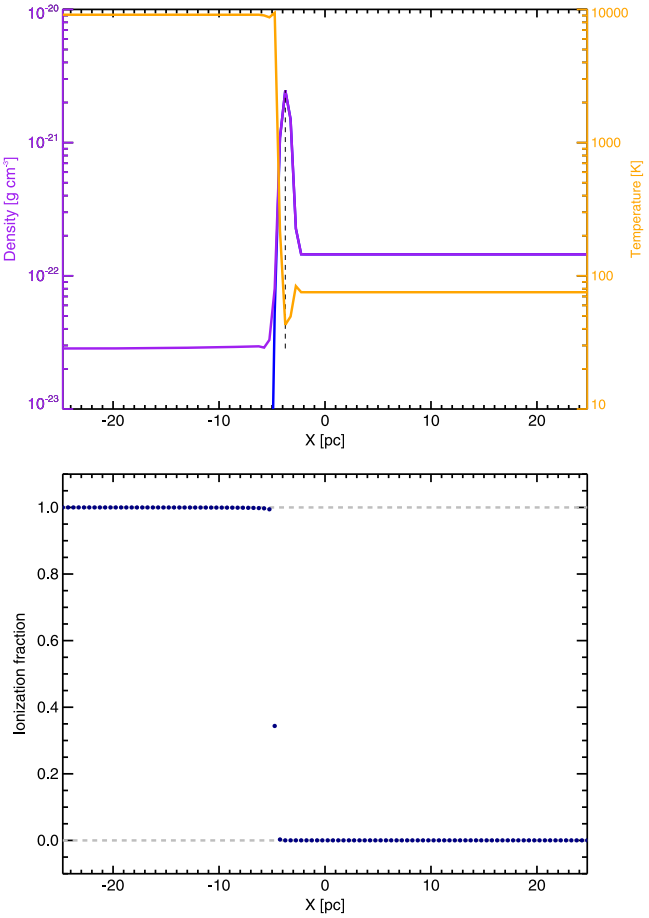

Figure 13. Ionization front structure at $t = 3.8$ Myr, with total density (purple), neutral density (blue), and temperature (orange), shown in the upper plot. The lower plot shows the ionization fraction. The dashed line in the upper plot is the front location, used for Figure 14. To the right of the front location, the neutral density is equal to the total density and is thus hidden.

as the peak in density, separating the lower density ionized region from the higher density neutral background.

Quantitatively, we can compare our simulation output to analytic calculations of ionization front expansion in H II regions. When an ionization source first turns on, the initial rate of photoionization overwhelms the rate of recombination. The result is a rapid phase of ionization and expansion, faster than the local sound speed. As this R-type expansion slows down due to reaching near ionization equilibrium at the Strömgren length, D-type expansion occurs, driven by an overpressure in the photoionized region. At times much greater than the sound crossing time of the ionized region, the photoionized gas will reach uniform pressure and density and exert a pressure on the swept up front of neutral material that bounds the ionized region. We can use the requirement from momentum conservation, that the photoionized region's force on the neutral gas be equal to the neutral gas's rate of momentum change per unit area, to derive the size of the ionized region. The ionized region's pressure on the swept up, neutral gas is

$$P_{\rm i} = c_{\rm s,i}^2 \rho_{\rm i} = c_{\rm s,i}^2 \mu_{\rm i} n_{\rm e} = c_{\rm s,i}^2 \mu_{\rm i} \sqrt{\frac{F_0}{\alpha_B l}}, \qquad (27)$$

where l is the size of the photoionized region, $c_{\rm s,i}$ is the sound speed in the photoionized region, and $\mu_{\rm i} = \mu_{\rm H}/2$ is the mean

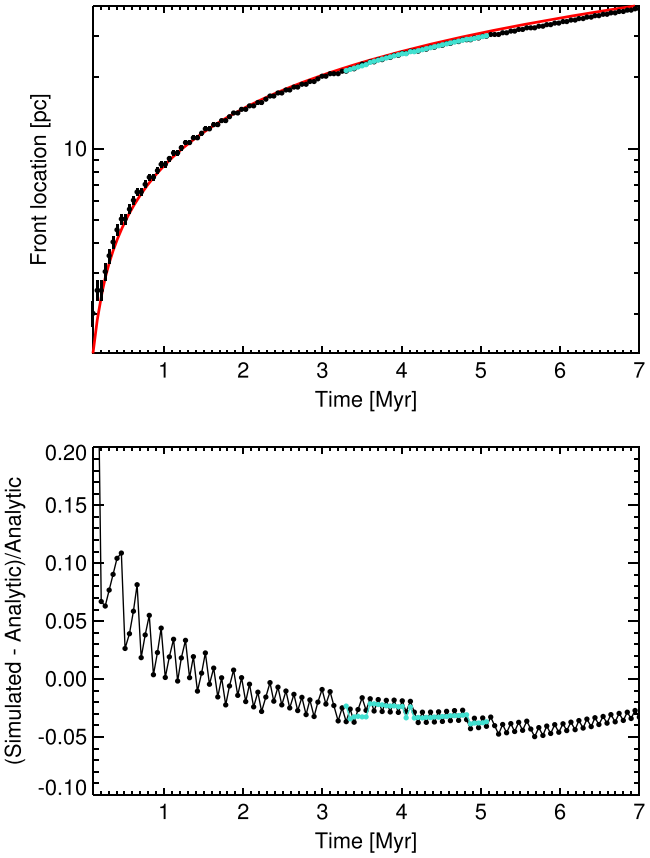

Figure 14. Comparison of the ionization front location for a coarse domain (black), with an SMR region (blue), to D-type analytic approximation (red) as a function of time (top), with residuals (bottom).

gas mass of the ions. We assume that the electron number density $n_{\rm e}$ equals the ion number density $n_{\rm i}$ because the photoionized region is in ionization equilibrium. At late times, the gas will be in thermal equilibrium, so we assume that $c_{\rm s,i}$ is a constant. The time rate of change of the neutral material's momentum is

$$P_0 = \frac{1}{A}\frac{d}{dt}(mv) = \frac{1}{A}\frac{d}{dt}\left(n_0 \mu A \Delta l \dot{l}\right) \approx \frac{d}{dt}\left(n_0 \mu l \dot{l}\right), \qquad (28)$$

where A is the cross-sectional area of the front and we have assumed that the distance over which mass has been swept, Δl, is comparable to the size l of the ionized region. Equating Equations (27) and (28) and assuming a similarity solution of the form $l \propto t^n$ yields

$$l(t) = C' t^{4/5}, \quad C' = \left(\frac{25}{12}\sqrt{\frac{F}{\alpha_B}}\frac{c_{\rm s,i}^2 \mu_{\rm i}}{n_0 \mu}\right)^{2/5}. \qquad (29)$$

Figure 14 shows the agreement between our simulation and analytic approximation for a coarse (black) and finer SMR region (blue). For this plot, the front location, or size of the ionized region corresponds to the x-location of maximum density, for a slice through the box at fixed y and z; an example of this location is highlighted by the dashed line in Figure 13. Choice of other slices orthogonal to the radiation direction yields similar results, as the gas is symmetric in both of these directions. We determine the corresponding analytic front location, using the simulated ionized sound speed as input into

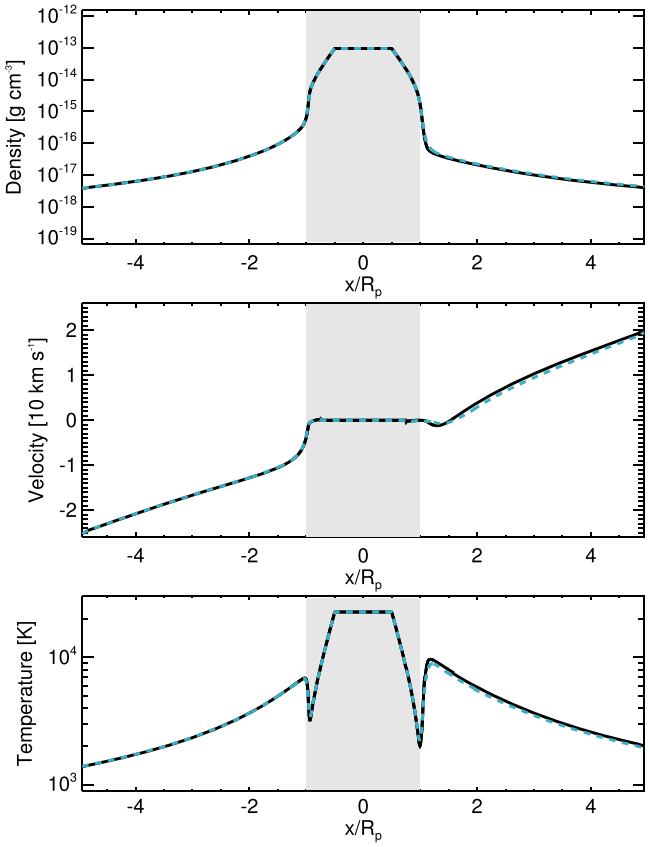

Figure 15. High resolution (solid) and lower resolution (dashed blue) simulations have well-converged outflow results, as shown in these 1D profiles.

Equation (29). The result we find is that the front location agrees with the analytic expression to within 5%. The oscillatory behavior of the residuals is a result of output that is not commensurate with the propagation time across one cell. Nevertheless, we find that the coarse and fine regions have consistent agreement with the analytic approximation. While we have tested the physical validity of our radiative transfer algorithm with SMR and MPI, we have not examined its efficiency and scaling.

APPENDIX B
CONVERGENCE OF MASS LOSS SIMULATIONS

To validate our numerical results, we run a convergence test by decreasing the resolution of the planet's atmosphere by a factor of two. Since additional refinement levels increase the resolution of our planet, not the entire box, we re-run our fiducial simulation with four levels of refinement, instead of five.

As highlighted in the 1D profiles in Figure 15, these two simulations are converged and produce comparable outflows. As expected, differences arise only in the stability of the static atmosphere, since lower resolution decreased the number of resolved scale lengths. Thus, the only differences in this resolution test were seen around the edge of the atmosphere, not in the escaping gas. We find no difference in the calculated mass loss rates for the two resolutions, further suggesting that our results are converged.

In the main text, we present our high resolution results for our fiducial parameters. Due to the convergence we describe here, from our resolution tests of four and five levels of refinement, we use only four levels of refinement for the low-flux simulation described in the main text.

REFERENCES

Adams, F. C. 2011, ApJ, 730, 27
Ballester, G. E., & Ben-Jaffel, L. 2015, ApJ, 804, 116
Baraffe, I., Selsis, F., Chabrier, G., et al. 2004, A&A, 419, L13
Ben-Jaffel, L. 2008, ApJ, 688, 1352
Bento, J., Wheatley, P. J., Copperwheat, C. M., et al. 2014, MNRAS, 437, 1511
Bisikalo, D., Kaygorodov, P., Ionov, D., et al. 2013, ApJ, 764, 19
Black, J. H. 1981, MNRAS, 197, 553
Bourrier, V., & Lecavelier des Etangs, A. 2013, A&A, 557, A124
Chamberlain, J. W. 1963, P&SS, 11, 901
Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor, M. 2000, ApJL, 529, L45
Childs, H., Brugger, E., Whitlock, B., et al. 2012, in High Performance Visualization: Enabling Extreme-scale Scientific Insight, ed. E. W. Bethel (Boca Raton, FL: CRC Press), 357
Cohen, O., Kashyap, V. L., Drake, J. J., et al. 2011, ApJ, 733, 67
Ehrenreich, D., Bourrier, V., Wheatley, P. J., et al. 2015, Natur, 522, 459
Ehrenreich, D., Lecavelier des Etangs, A., Hébrard, G., et al. 2008, A&A, 483, 933
García Muñoz, A. 2007, P&SS, 55, 1426
Gendelev, L., & Krumholz, M. R. 2012, ApJ, 745, 158
Gross, S. H. 1972, JAtS, 29, 214
Haswell, C. A., Fossati, L., Ayres, T., et al. 2012, ApJ, 760, 79
Henry, G. W., Marcy, G. W., Butler, R. P., & Vogt, S. S. 2000, ApJL, 529, L41
Holmström, M., Ekenbäck, A., Selsis, F., et al. 2008, Natur, 451, 970
Koskinen, T. T., Harris, M. J., Yelle, R. V., & Lavvas, P. 2013, Icar, 226, 1678
Koskinen, T. T., Lavvas, P., Harris, M. J., & Yelle, R. V. 2014, RSPTA, 372, 30089
Krumholz, M. R., Stone, J. M., & Gardiner, T. A. 2007, ApJ, 671, 518
Kulow, J. R., France, K., Linsky, J., & Loyd, R. O. P. 2014, ApJ, 786, 132
Lammer, H., Selsis, F., Ribas, I., et al. 2003, ApJL, 598, L121
Lecavelier des Etangs, A., Bourrier, V., Wheatley, P. J., et al. 2012, A&A, 543, L4
Lecavelier des Etangs, A., Ehrenreich, D., Vidal-Madjar, A., et al. 2010, A&A, 514, A72
Li, S.-L., Miller, N., Lin, D. N. C., & Fortney, J. J. 2010, Natur, 463, 1054
Linsky, J. L., Yang, H., France, K., et al. 2010, ApJ, 717, 1291
Llama, J., & Shkolnik, E. L. 2015, ApJ, 802, 41
Llama, J., Vidotto, A. A., Jardine, M., et al. 2013, MNRAS, 436, 2179
Lopez, E. D., & Fortney, J. J. 2013, ApJ, 776, 2
Matsakos, T., Uribe, A., & Königl, A. 2015, A&A, 578, A6
Menager, H., Barthélemy, M., Koskinen, T., et al. 2013, Icar, 226, 1709
Morton, D. C. 2003, ApJS, 149, 205
Murray-Clay, R. A., Chiang, E. I., & Murray, N. 2009, ApJ, 693, 23
Osterbrock, D. E. 1989, Astrophysics of Gaseous Nebulae and Active Galactic Nuclei (Mill Valley, CA: Univ. Science Books)
Owen, J. E., & Adams, F. C. 2014, MNRAS, 444, 3761
Owen, J. E., & Jackson, A. P. 2012, MNRAS, 425, 2931
Owen, J. E., & Wu, Y. 2013, ApJ, 775, 105
Parker, E. N. 1958, ApJ, 128, 664
Poppenhaeger, K., Schmitt, J. H. M. M., & Wolk, S. J. 2013, ApJ, 773, 62
Quirk, J. J. 1994, IJNMF, 18, 555
Rappaport, S., Barclay, T., DeVore, J., et al. 2014, ApJ, 784, 40
Schneiter, E. M., Velázquez, P. F., Esquivel, A., Raga, A. C., & Blanco-Cano, X. 2007, ApJL, 671, L57
Sofia, U. J., & Meyer, D. M. 2001, ApJL, 554, L221
Southworth, J., Hinse, T. C., Dominik, M., et al. 2012, MNRAS, 426, 1338
Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, ApJS, 178, 137
Stone, J. M., & Proga, D. 2009, ApJ, 694, 205
Tenorio-Tagle, G., Bodenheimer, P., Lin, D. N. C., & Noriega-Crespo, A. 1986, MNRAS, 221, 635
Tian, F., Toon, O. B., Pavlov, A. A., & De Sterck, H. 2005, ApJ, 621, 1049
Tóth, G., & Roe, P. L. 2002, JCoPh, 180, 736
Trammell, G. B., Arras, P., & Li, Z.-Y. 2011, ApJ, 728, 152
Trammell, G. B., Li, Z.-Y., & Arras, P. 2014, ApJ, 788, 161

Tremblin, P., & Chiang, E. 2013, MNRAS, 428, 2565
Vidal-Madjar, A., Désert, J.-M., Lecavelier des Etangs, A., et al. 2004, ApJL, 604, L69
Vidal-Madjar, A., Lecavelier des Etangs, A., Désert, J.-M., et al. 2003, Natur, 422, 143
Watson, A. J., Donahue, T. M., & Walker, J. C. G. 1981, Icar, 48, 150
Whalen, D., Abel, T., & Norman, M. L. 2004, ApJ, 610, 14
Woods, T. N., Rottman, G. J., Bailey, S. M., Solomon, S. C., & Worden, J. R. 1998, SoPh, 177, 133
Yelle, R. V. 2004, Icar, 170, 167