A. [This problem is not treatable by acceptance-rejection sampling, because the support of Y is unbounded.] The CDF of Y is

$$F_Y(y) = \int y e^{-y^2/2} \, dy = -e^{-y^2/2} + C.$$

Because the CDF must limit to 1 as $y \to \infty$, we conclude that $C = 1$ and $F_Y(y) = 1 - e^{-y^2/2}$. Setting this expression equal to x and solving for y, we compute

$$\begin{aligned} x &= 1 - e^{-y^2/2} \\ \Rightarrow \quad e^{-y^2/2} &= 1 - x \\ \Rightarrow \quad -y^2/2 &= \log(1-x) \\ \Rightarrow \quad y^2 &= -2\log(1-x) \\ \Rightarrow \quad y &= \sqrt{-2\log(1-x)}. \end{aligned}$$

Thus $F_Y^{-1}(x) = \sqrt{-2\log(1-x)}$ for $0 < x < 1$. The inverse transform algorithm for generating a random value y of Y is:

1. Choose a number x uniformly randomly on $[0, 1]$.
2. Compute $y = \sqrt{-2\log(1-x)}$.

B. For starters,

$$F_T(t) = P(T \leq t) = P(Y^X \leq t).$$

The probability on the right equals the integral of the joint density $f_{X,Y}(x, y)$ over a certain region of the x-y-plane. The region is bounded by the line $x = 0$, the line $y = 0$, and the curve $y = t^{1/x}$. [A sketch is helpful.] The integral is

$$\int_0^\infty \int_0^{t^{1/x}} f_{X,Y}(x, y) \, dy \, dx.$$

Meanwhile, the joint density splits as $f_{X,Y}(x, y) = f_X(x) f_Y(y)$ because X and Y are independent. Therefore the integral is

$$\int_0^\infty \int_0^{t^{1/x}} f_X(x) f_Y(y) \, dy \, dx = \int_0^\infty f_X(x) F_Y\left(t^{1/x}\right) \, dx.$$

Finally,

$$\begin{aligned}
f_T(t) &= \frac{d}{dt} F_T(t) \\
&= \frac{d}{dt} \int_0^\infty f_X(x) F_Y\left(t^{1/x}\right)\,dx \\
&= \int_0^\infty f_X(x) \frac{d}{dt} F_Y\left(t^{1/x}\right)\,dx \\
&= \int_0^\infty f_X(x) f_Y\left(t^{1/x}\right) \frac{d}{dt}\left(t^{1/x}\right)\,dx \\
&= \int_0^\infty f_X(x) f_Y\left(t^{1/x}\right) \frac{1}{x} t^{1/x-1}\,dx.
\end{aligned}$$

[We do not know the support of T, although we know that it is contained in $(0, \infty)$. Also, computing the integral in the $dx\ dy$ order is more difficult. The cases $t > 1$, $= 1$, and $t < 1$ should be treated separately, and the integral for $t > 1$ should be written as a sum of two integrals. In the end, I get

$$f_T(t) = \int_0^\infty f_Y(y) f_X\left(\frac{\log t}{\log y}\right) \cdot \frac{1}{t \log y}\,dy$$

for $t > 1$ and

$$f_T(t) = \int_0^1 f_Y(y) f_X\left(\frac{\log t}{\log y}\right) \cdot \frac{-1}{t \log y}\,dy$$

for $t < 1$. The $dy\ dx$ treatment suggests that these two expressions should agree as $t \to 1$. That is not obvious.]

C.A. We compute

$$\begin{aligned}
m_X(t) &= E\left(e^{tX}\right) \\
&= \sum_{k=0}^\infty e^{tk} P(X = k) \\
&= \sum_{k=0}^\infty e^{tk} (1-p)^k p \\
&= p \sum_{k=0}^\infty \left(e^t (1-p)\right)^k \\
&= \frac{p}{1 - e^t(1-p)}
\end{aligned}$$

for t sufficiently close to 0.

C.B. Recall that Y can be regarded as the sum of r independent Geom(p) random variables. Therefore

$$m_Y(t) = (m_X(t))^r = \left(\frac{p}{1 - e^t(1-p)}\right)^r.$$

C.C. Because Y is a sum of r IID random variables (with finite mean and variance), the central limit theorem says that Y is approximately normal as $r \to \infty$. But $e^{t^2/2}$ is the MGF of a *standard* normal random variable. So to make $m_Z(t) \to e^{t^2/2}$ we should let Z be the standardization of Y. More explicitly, we know that $E(Y) = r(1-p)/p$ and $V(Y) = r(1-p)/p^2$. So

$$Z = \frac{Y - E(Y)}{SD(Y)} = \frac{Y - r(1-p)/p}{\sqrt{r(1-p)/p^2}}.$$

D.A. [This problem is Theorem 9.3.2, which was in the assigned reading but not discussed in class. In my opinion, this problem is the most difficult one on the exam.] Consider

$$E(Y|X = x) \cdot h(x),$$

which is a function of x. Because x and $h(x)$ are not random variables (but rather fixed values of X and $h(X)$, respectively), we can apply linearity of expectation:

$$E(Y|X = x) \cdot h(x) = E(Y \cdot h(x)|X = x).$$

When we turn this function of x into a function of X, we obtain

$$E(Y|X) \cdot h(X) = E(Y \cdot h(X)|X).$$

D.B. [This problem is Theorem 9.3.9, which was in the assigned reading but not discussed in class.] By the definition of covariance,

$$\text{Cov}(Y - E(Y|X), h(X)) = E((Y - E(Y|X))h(X)) - E(Y - E(Y|X)) \cdot E(h(X)).$$

We now show that both terms on the right side of this equation are zero. By linearity of expectation and the law of total expectation, the right term is

$$(E(Y) - E(E(Y|X))) \cdot E(h(X)) = 0 \cdot E(h(X)) = 0.$$

By simple algebra and Problem D.A, the left term is

$$E(Yh(X) - E(Y|X)h(X)) = E(Yh(X) - E(Yh(X)|X)).$$

Then, by linearity of expectation and the law of total expectation, the left term becomes

$$E(Yh(X)) - E(E(Yh(X)|X)) = 0.$$

Therefore the covariance of $Y - E(Y|X)$ and $h(X)$ is zero, and they are uncorrelated.

E.A. How many visitors arrive today? The answer is $X \sim \text{Pois}(1{,}000)$, if we model arrivals as a Poisson process over one day of time.

E.B. How much time passes between consecutive arrivals? The answer is $Y \sim \text{Expo}(1{,}000)$, if we model arrivals as a Poisson process over one day of time.

E.C. What is the average number of visitors over n days? It's S_n/n, where $S_n = X_1 + \cdots + X_n$ and $X_i \sim \text{Pois}(1{,}000)$ is the number of visitors on the ith day. Then

$$E(S_n/n) = \frac{1}{n} nE(X) = \lambda$$

and

$$V(S_n/n) = \frac{1}{n^2} nV(X) = \lambda/n$$

because the X_i are independent. By the central limit theorem, S_n/n is approximately distributed as $\text{Norm}(\lambda, \lambda/n)$.

E.D. What is the quantile of today's visitor count, among the visitor counts of all days? It's $U \sim \text{Unif}(0, 1)$. [See Exam B Problem B.]

E.E. Bernoulli: How many current presidents of the USA are visiting my site today? The answer is either 0 or 1, and it's random, so it's $T \sim \text{Bern}(p)$ for some p. To get an idea of how p relates to $\lambda = 1{,}000$, assume that all $m = 8{,}000{,}000{,}000$ people in the world are equally probable to visit my site today and that their visits are independent of each other. Then the number of visitors to my site today is $V \sim \text{Binom}(m, p)$. Combining our two models, we expect the number of visitors today to be

$$mp = E(V) = E(X) = \lambda.$$

Thus

$$p = \lambda/m = \frac{1}{8{,}000{,}000}.$$

[Here's another solution that would also earn full credit, inspired by the binomial-Poisson relationship discussed in class: How many visitors do I receive in the next millisecond? It is quite improbable that I receive more than one visitor in the next millisecond. Therefore the question is answered approximately by $T \sim \text{Bern}(p)$ for some p. Now what is p? We'll figure it out from the expected number of visitors in the next millisecond. The Bernoulli treatment says that it's approximately $E(T) = p$. A Poisson treatment says that it's $1{,}000/(24 \cdot 60 \cdot 60 \cdot 1{,}000)$. So $p \approx 1/(24 \cdot 60 \cdot 60)$.]