A5 (by Abhinav Kumar) Define $G:\mathbb{R}\to\mathbb{R}$ by $G(x) = \int_0^x g(t)\,dt$. By assumption, G is a strictly increasing, thrice continuously differentiable function. It is also bounded: for $x>1$, we have

$$0 < G(x) - G(1) = \int_1^x g(t)\,dt \leq \int_1^x dt/t^2 = 1,$$

and similarly, for $x<-1$, we have $0 > G(x) - G(-1) \geq -1$. It follows that the image of G is some open interval (A,B) and that $G^{-1}:(A,B)\to\mathbb{R}$ is also thrice continuously differentiable.

Define $H:(A,B)\times(A,B)\to\mathbb{R}$ by $H(x,y) = F(G^{-1}(x),G^{-1}(y))$; it is twice continuously differentiable since F and G^{-1} are. By our assumptions about F,

$$\begin{aligned}\frac{\partial H}{\partial x} + \frac{\partial H}{\partial y} &= \frac{\partial F}{\partial x}(G^{-1}(x),G^{-1}(y))\cdot\frac{1}{g(G^{-1}(x))} \\ &\quad + \frac{\partial F}{\partial y}(G^{-1}(x),G^{-1}(y))\cdot\frac{1}{g(G^{-1}(y))} = 0.\end{aligned}$$

Therefore H is constant along any line parallel to the vector $(1,1)$, or equivalently, $H(x,y)$ depends only on $x-y$. We may thus write $H(x,y) = h(x-y)$ for some function h on $(-(B-A),B-A)$, and we then have $F(x,y) = h(G(x)-G(y))$. Since $F(u,u)=0$, we have $h(0)=0$. Also, h is twice continuously differentiable (since it can be written as $h(x) = H((A+B+x)/2,(A+B-x)/2)$), so $|h'|$ is bounded on the closed interval $[-(B-A)/2,(B-A)/2]$, say by M.

Given $x_1,\dots,x_{n+1}\in\mathbb{R}$ for some $n\geq 2$, the numbers $G(x_1),\dots,G(x_{n+1})$ all belong to (A,B), so we can choose indices i and j so that $|G(x_i)-G(x_j)| \leq (B-A)/n \leq (B-A)/2$. By the mean value theorem,

$$|F(x_i,x_j)| = |h(G(x_i)-G(x_j))| \leq M\frac{B-A}{n},$$

so the claim holds with $C = M(B-A)$.

A6 Choose some ordering $h_1,\dots,h_n$ of the elements of G with $h_1 = e$. Define an $n\times n$ matrix M by settting $M_{ij} = 1/k$ if $h_j = h_i g$ for some $g\in\{g_1,\dots,g_k\}$ and $M_{ij}=0$ otherwise. Let v denote the column vector $(1,0,\dots,0)$. The probability that the product of m random elements of $\{g_1,\dots,g_k\}$ equals h_i can then be interpreted as the i-th component of the vector $M^m v$.

Let $\hat{G}$ denote the dual group of G, i.e., the group of complex-valued characters of G. Let $\hat{e}\in\hat{G}$ denote the trivial character. For each $\chi\in\hat{G}$, the vector $v_\chi = (\chi(h_i))_{i=1}^n$ is an eigenvector of M with eigenvalue $\lambda_\chi = (\chi(g_1)+\cdots+\chi(g_k))/k$. In particular, $v_{\hat{e}}$ is the all-ones vector and $\lambda_{\hat{e}} = 1$. Put

$$b = \max\{|\lambda_\chi| : \chi\in\hat{G}-\{\hat{e}\}\};$$

we show that $b\in(0,1)$ as follows. First suppose $b=0$; then

$$1 = \sum_{\chi\in\hat{G}}\lambda_\chi = \frac{1}{k}\sum_{i=1}^k\sum_{\chi\in\hat{G}}\chi(g_i) = \frac{n}{k}$$

because $\sum_{\chi\in\hat{(G)}}\chi(g_i)$ equals n for $i=1$ and 0 otherwise. However, this contradicts the hypothesis that $\{g_1,\dots,g_k\}$ is not all of G. Hence $b>0$. Next suppose $b=1$, and choose $\chi\in\hat{G}-\{\hat{e}\}$ with $|\lambda_\chi| = 1$. Since each of $\chi(g_1),\dots,\chi(g_k)$ is a complex number of norm 1, the triangle inequality forces them all to be equal. Since $\chi(g_1) = \chi(e) = 1$, χ must map each of $g_1,\dots,g_k$ to 1, but this is impossible because χ is a nontrivial character and $g_1,\dots,g_k$ form a set of generators of G. This contradiction yields $b<1$.

Since $v = \frac{1}{n}\sum_{\chi\in\hat{G}} v_\chi$ and $Mv_\chi = \lambda_\chi v_\chi$, we have

$$M^m v - \frac{1}{n}v_{\hat{e}} = \frac{1}{n}\sum_{\chi\in\hat{G}-\{\hat{e}\}}\lambda_\chi^m v_\chi.$$

Since the vectors v_χ are pairwise orthogonal, the limit we are interested in can be written as

$$\lim_{m\to\infty}\frac{1}{b^{2m}}(M^m v - \frac{1}{n}v_{\hat{e}})\cdot(M^m v - \frac{1}{n}v_{\hat{e}}).$$

and then rewritten as

$$\lim_{m\to\infty}\frac{1}{b^{2m}}\sum_{\chi\in\hat{G}-\{\hat{e}\}}|\lambda_\chi|^{2m} = \#\{\chi\in\hat{G} : |\lambda_\chi| = b\}.$$

By construction, this last quantity is nonzero and finite.

Remark. It is easy to see that the result fails if we do not assume $g_1 = e$: take $G = \mathbb{Z}/2\mathbb{Z}$, $n=1$, and $g_1 = 1$.

Remark. Harm Derksen points out that a similar argument applies even if G is not assumed to be abelian, provided that the operator $g_1+\cdots+g_k$ in the group algebra $\mathbb{Z}[G]$ is *normal*, i.e., it commutes with the operator $g_1^{-1}+\cdots+g_k^{-1}$. This includes the cases where the set $\{g_1,\dots,g_k\}$ is closed under taking inverses and where it is a union of conjugacy classes (which in turn includes the case of G abelian).

Remark. The matrix M used above has nonnegative entries with row sums equal to 1 (i.e., it corresponds to a Markov chain), and there exists a positive integer m such that M^m has positive entries. For any such matrix, the Perron-Frobenius theorem implies that the sequence of vectors $M^m v$ converges to a limit w, and there exists $b\in[0,1)$ such that

$$\limsup_{m\to\infty}\frac{1}{b^{2m}}\sum_{i=1}^n((M^m v - w)_i)^2$$

is nonzero and finite. (The intended interpretation in case $b=0$ is that $M^m v = w$ for all large m.) However, the limit need not exist in general.